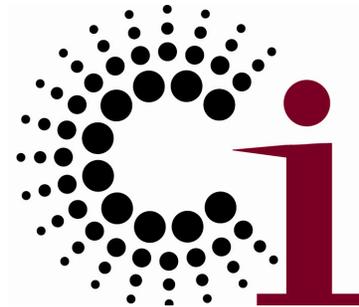




Looking to the Future

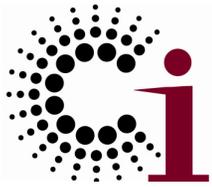
Ian Foster



Computation Institute

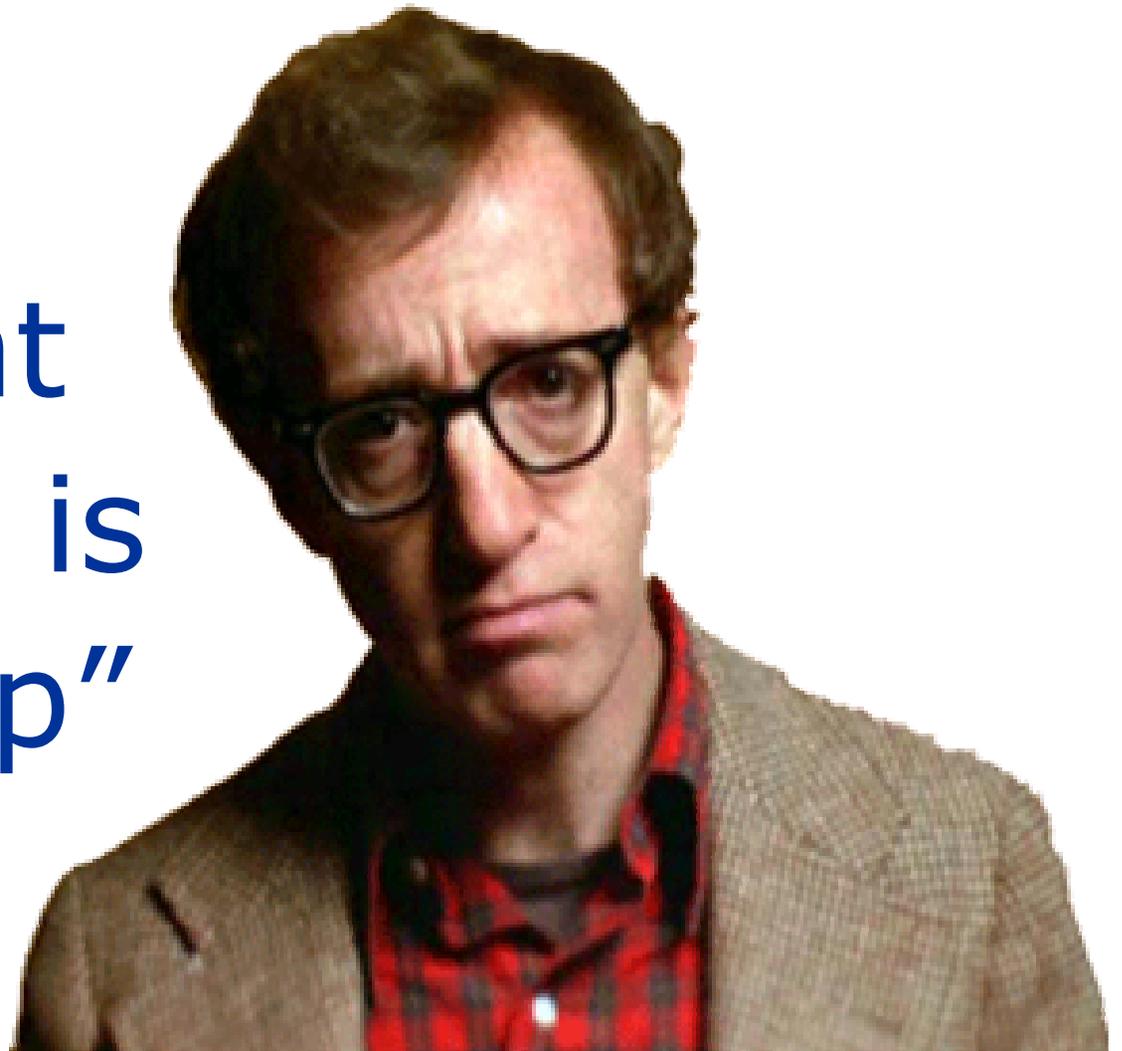
Argonne National Lab & University of Chicago

<http://ianfoster.typepad.com>

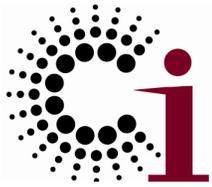


An American Philosopher Speaks on the Value of KAREN

“80 percent
of success is
showing up”



Woody Allen



J.C.R. Licklider Reflecting in 1960 on Where His Time Went



About 85 per cent
of my “thinking”
time was spent
getting into a
position to think,
to make a decision,
to learn something
I needed to know



For Example ...

“At one point, it was necessary to compare six experimental determinations of a function relating speech-intelligibility to speech-to-noise ratio.

No two experimenters had used the same definition or measure of speech-to-noise ratio. Several hours of

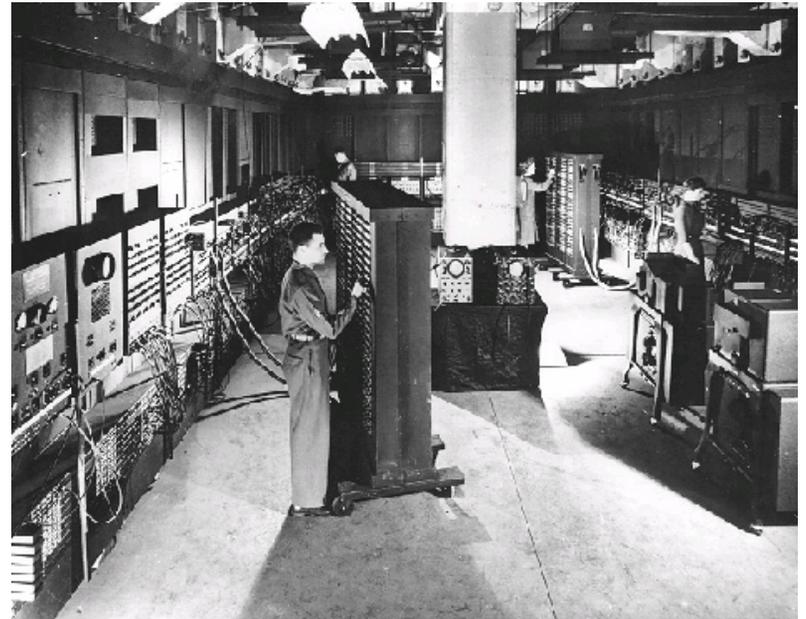
calculating were required to get the data into comparable form. When they were in comparable form, it took only a few seconds to determine what I needed to know.”





Man-Computer Symbiosis

“... will involve very close coupling between the human and the electronic members of the partnership ... Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much more effectively than man alone can perform them.”





Licklider's Vision

- He imagined
 - ◆ Low-cost personal computers
 - ◆ Organized, easily searchable data stores
 - ◆ Intuitive user interfaces
 - ◆ Computer aided reasoning
 - ◆ Interactive numerical simulation
- He didn't mention (at least in CMS paper)
 - ◆ Standard representations of knowledge
 - ◆ Collaborative problem solving among people

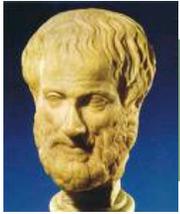


Fast Forward to 2007

“Compare six experimental determinations of a function relating speech-intelligibility to speech-to-noise ratio”



Emergence of New Problem Solving Methodologies



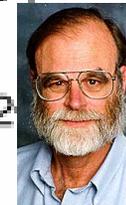
Empirical



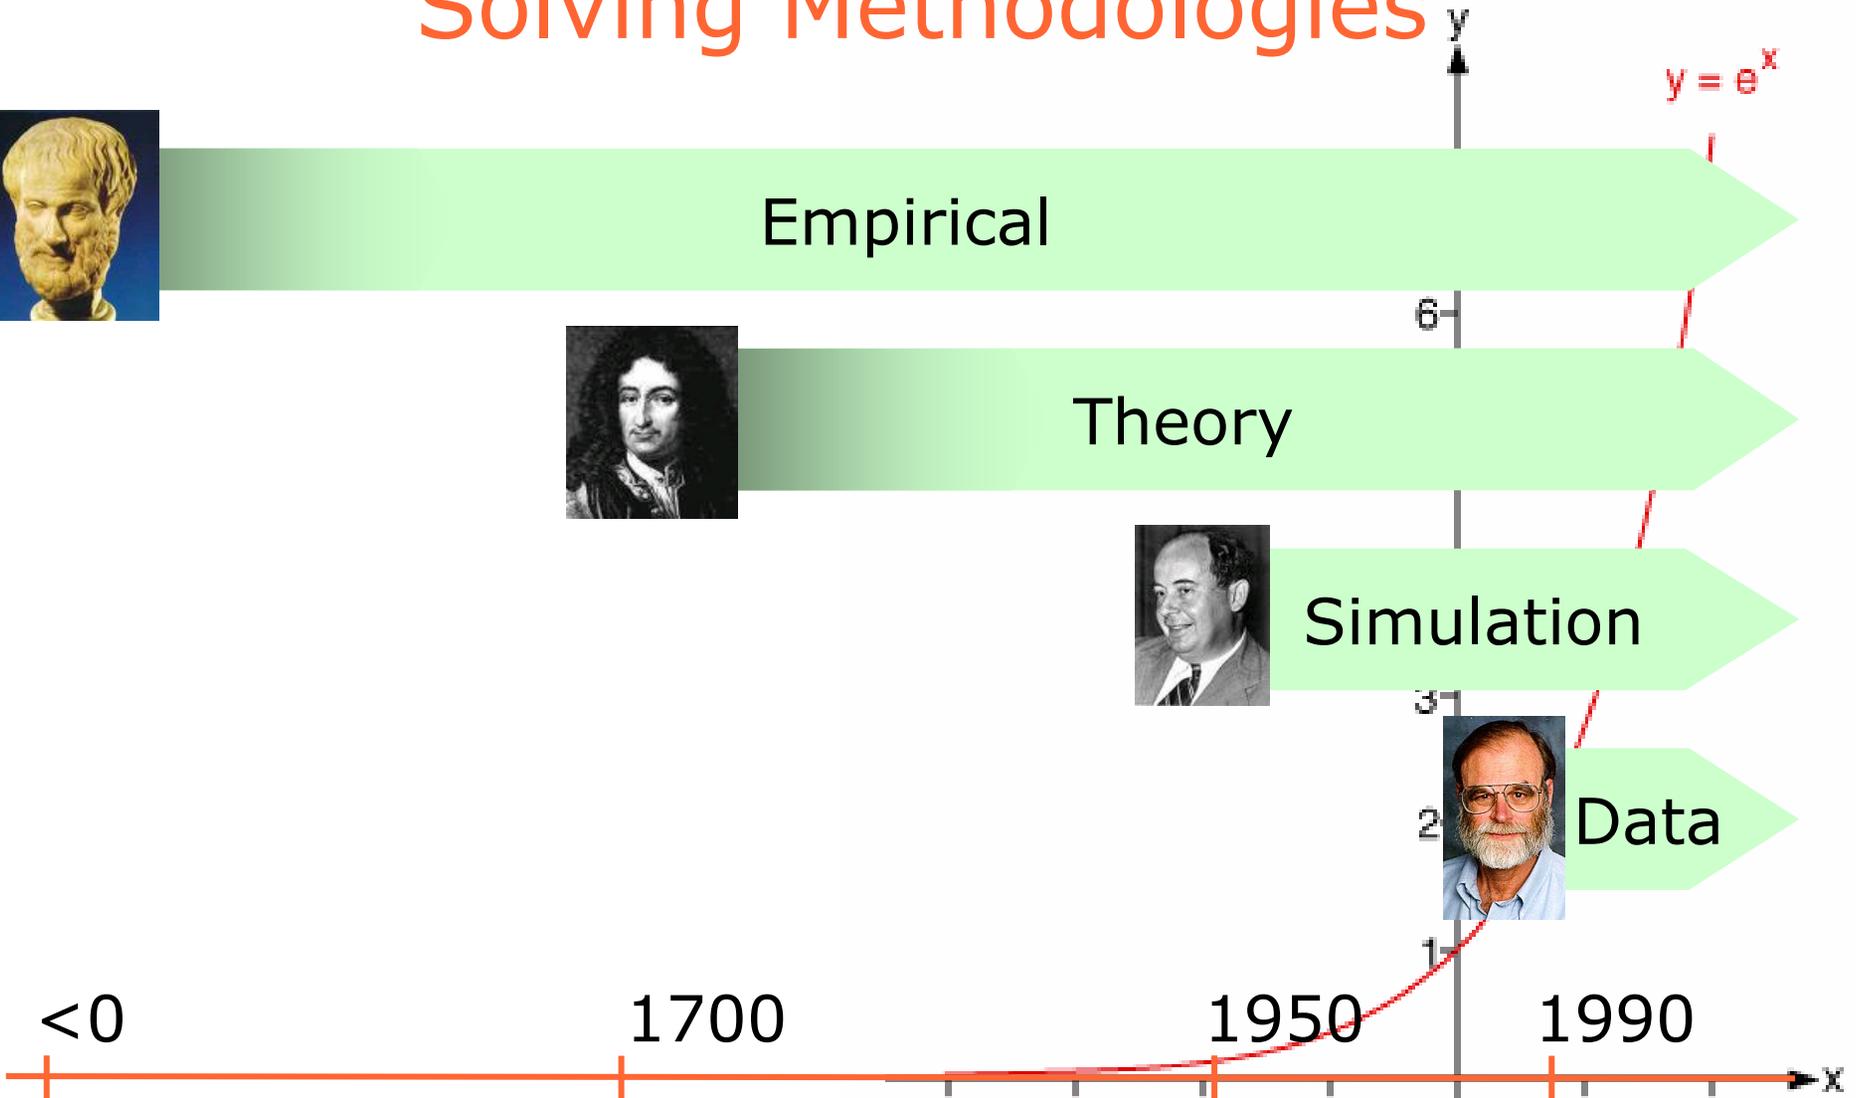
Theory



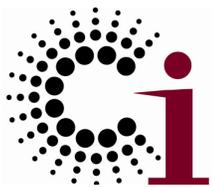
Simulation



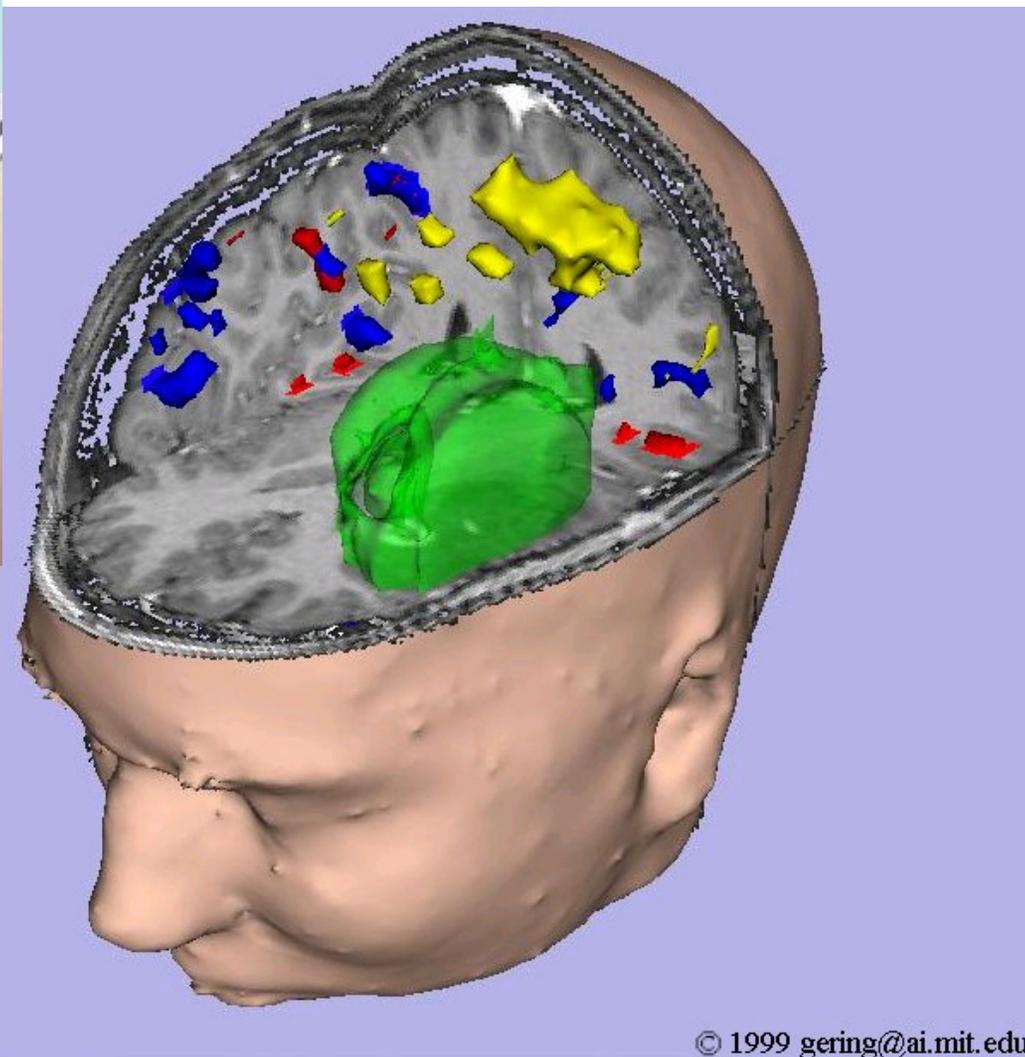
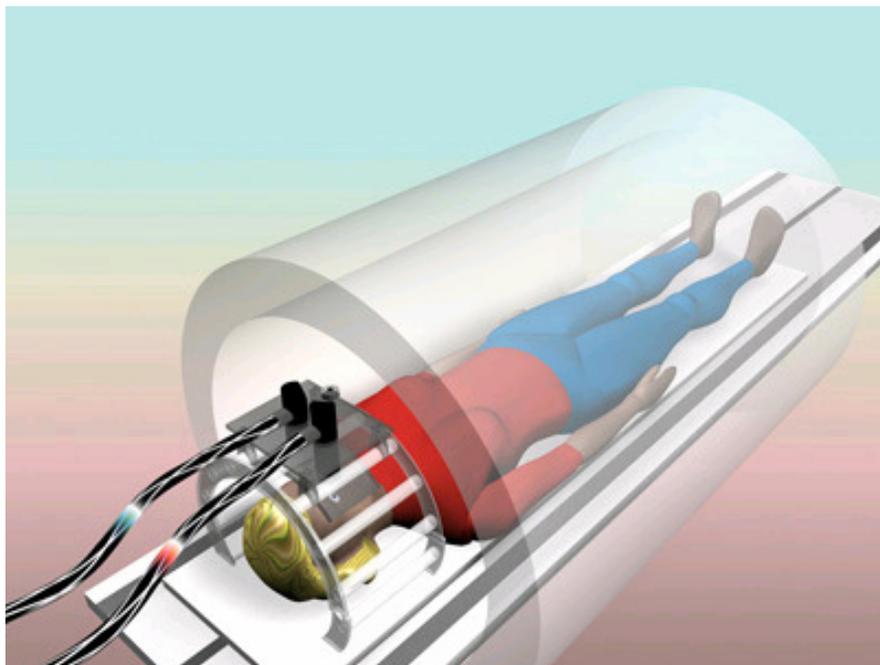
Data

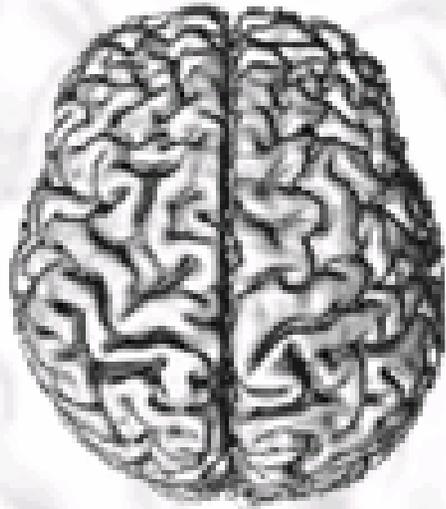


eScience: When brute force doesn't work anymore (Szalay)



Functional Magnetic Resonance Imaging (fMRI)





A public repository of peer-reviewed fMRI studies and their underlying data.

The fMRI Data
Center
(www.fmridc.org)



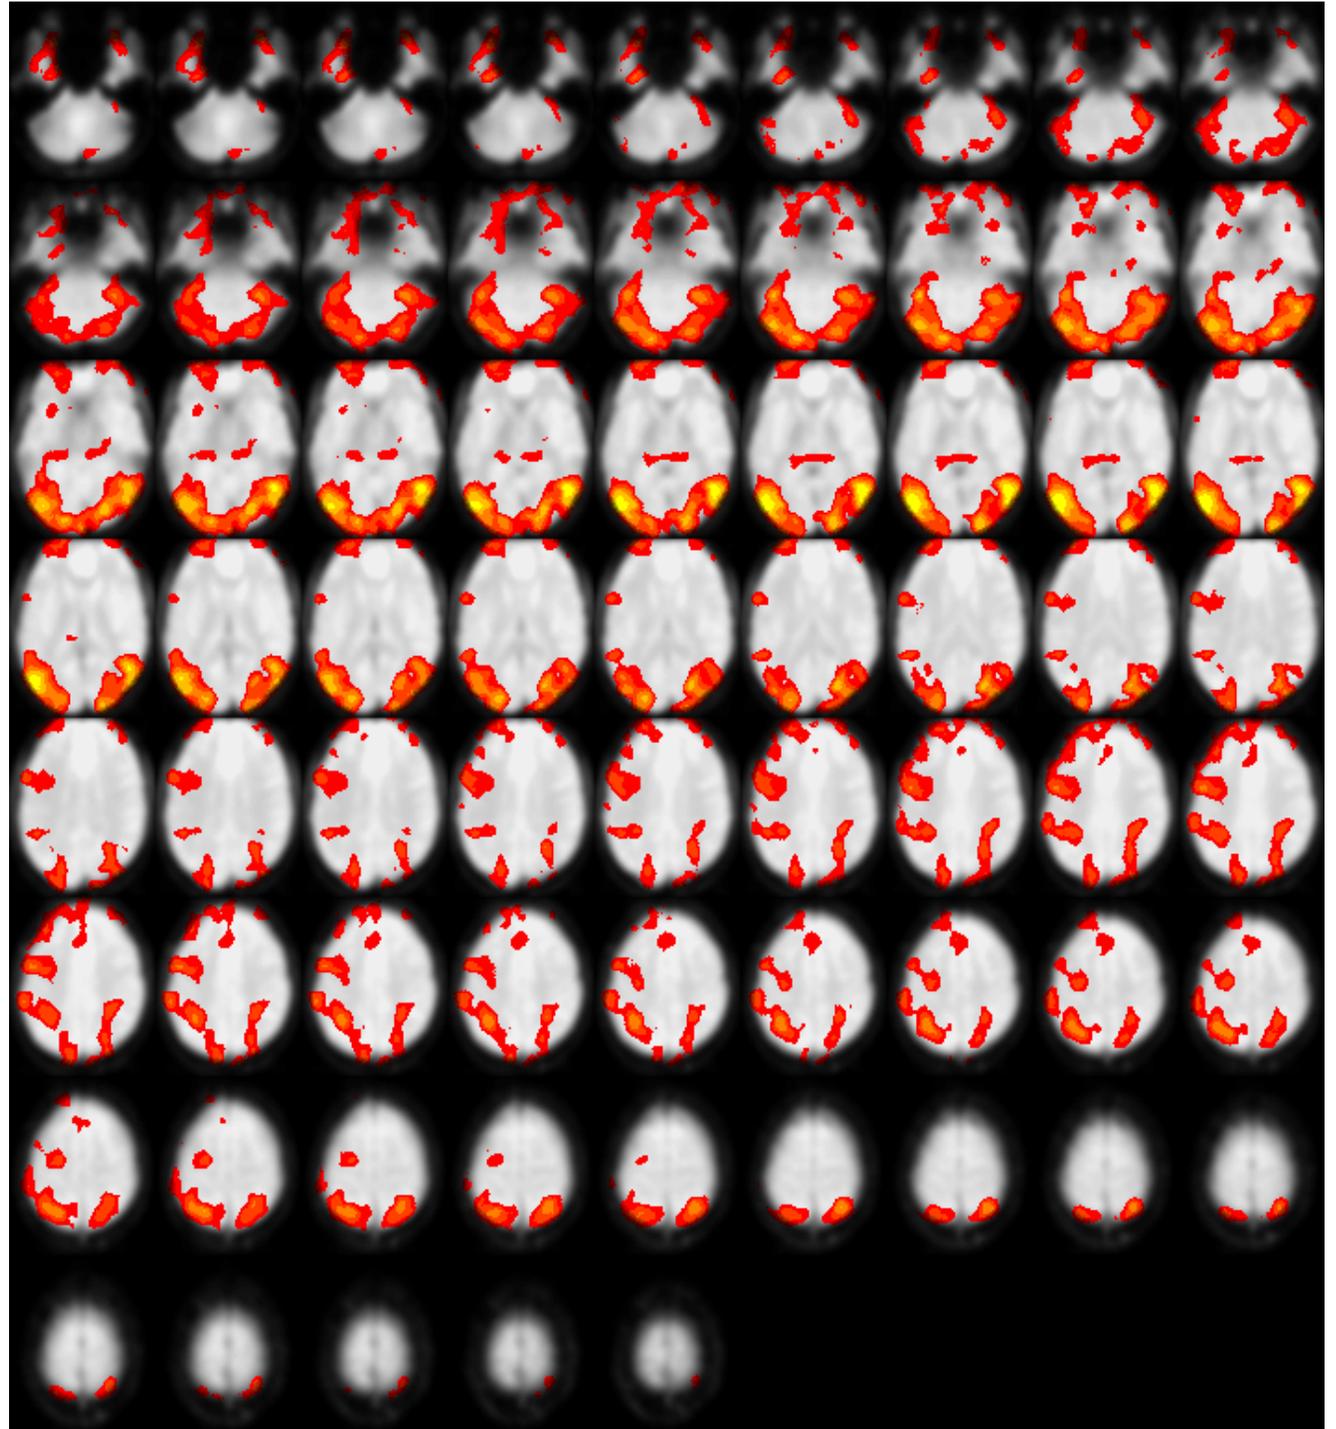
- A typical study comprises
 - 3 groups,
 - 20 subjects/group,
 - 5 runs/subject,
 - 300 volumes/run

→ 90,000 volumes, 60 GB raw

→ 1.2 million files processed
- 100s of such studies in total

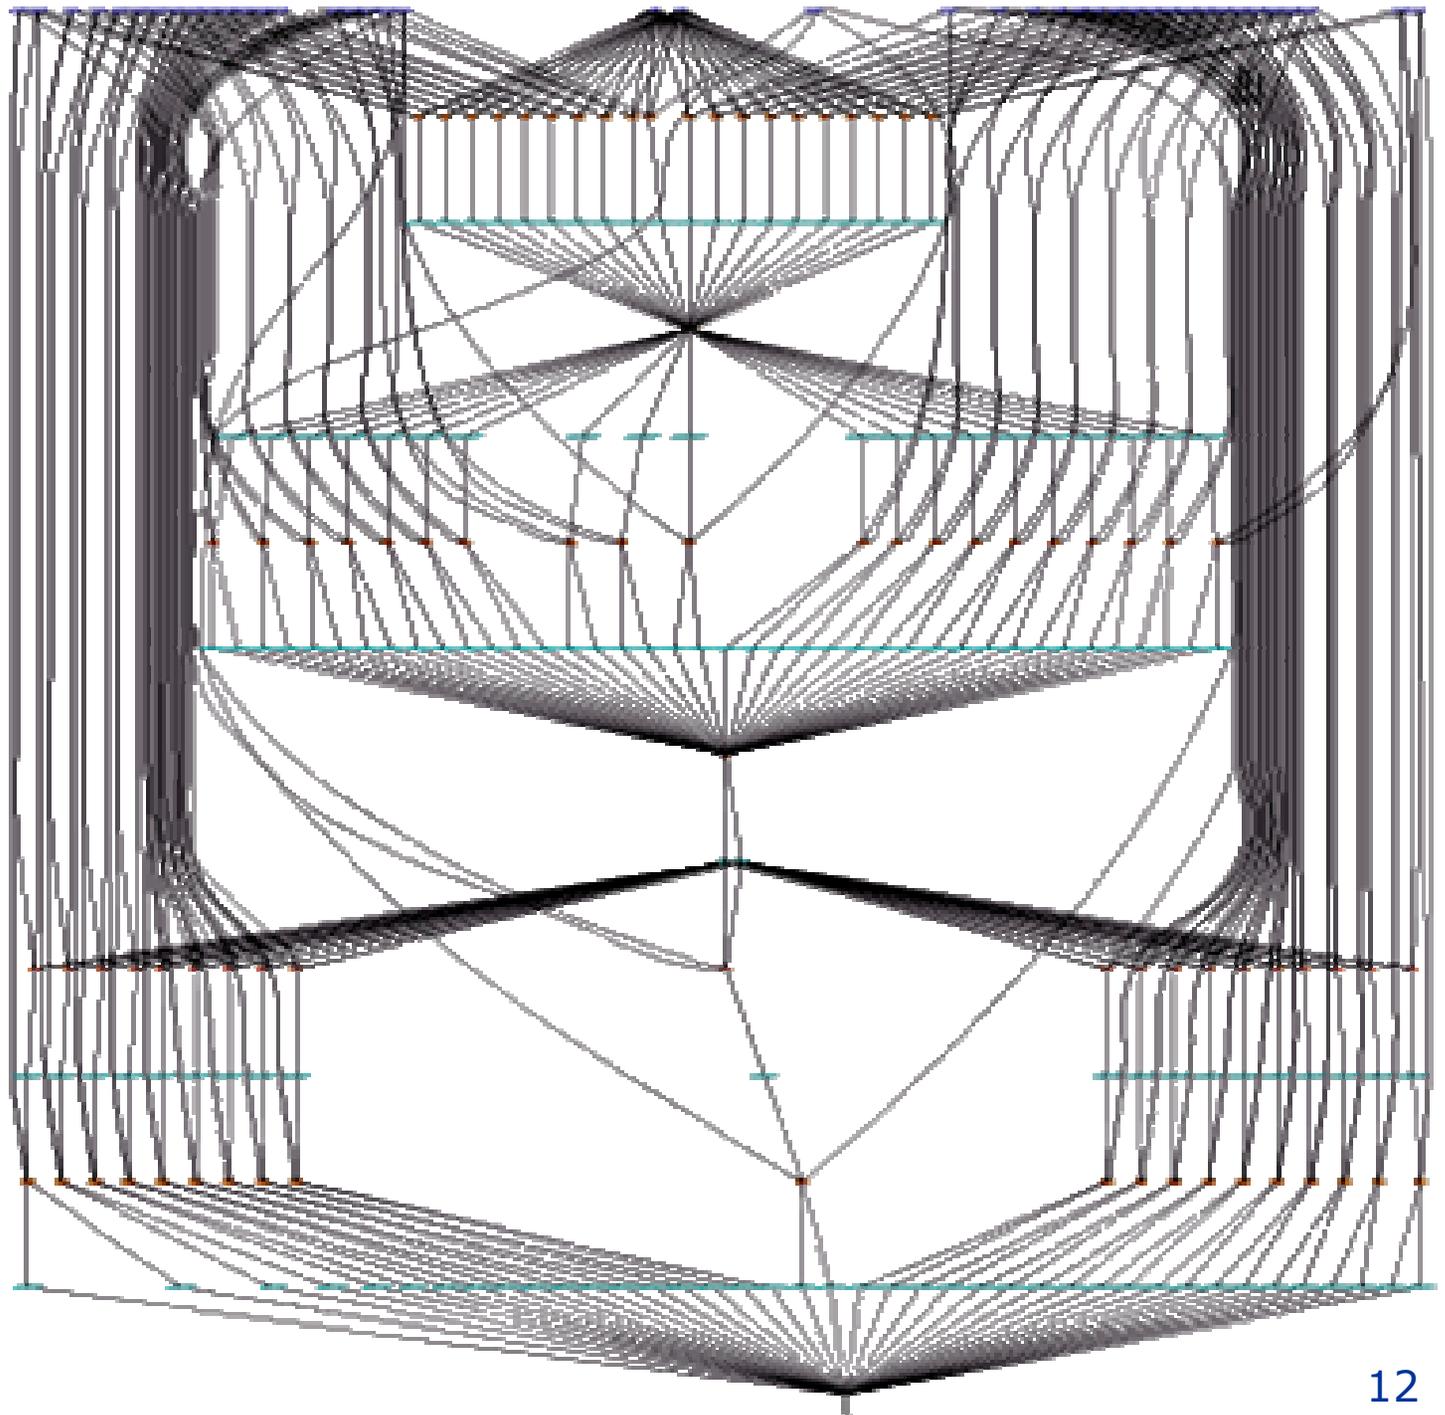


Functional MRI Analysis





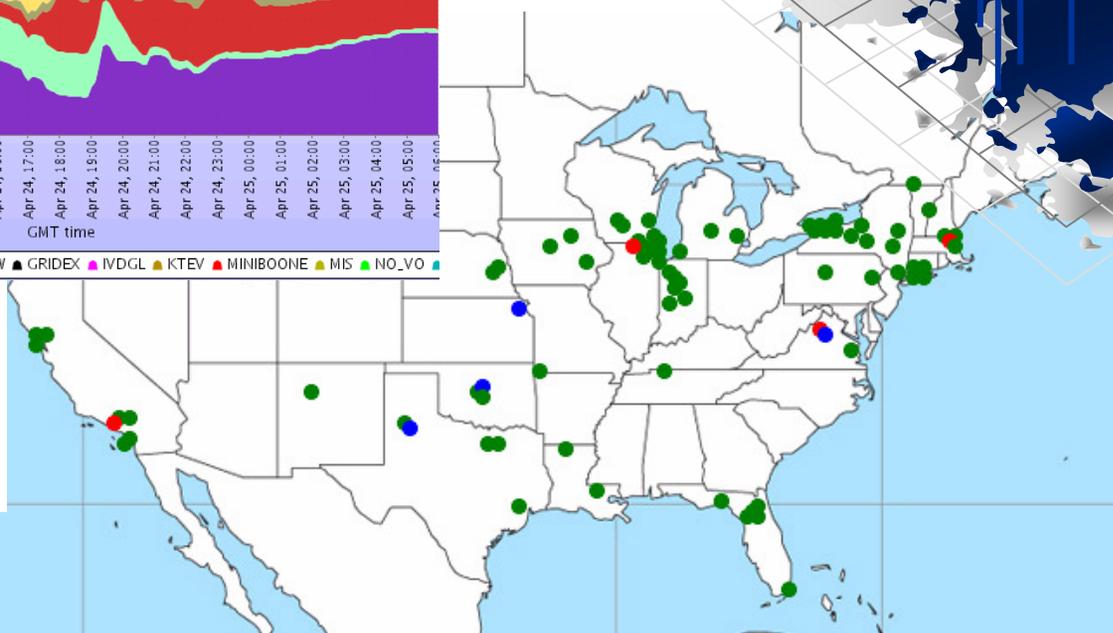
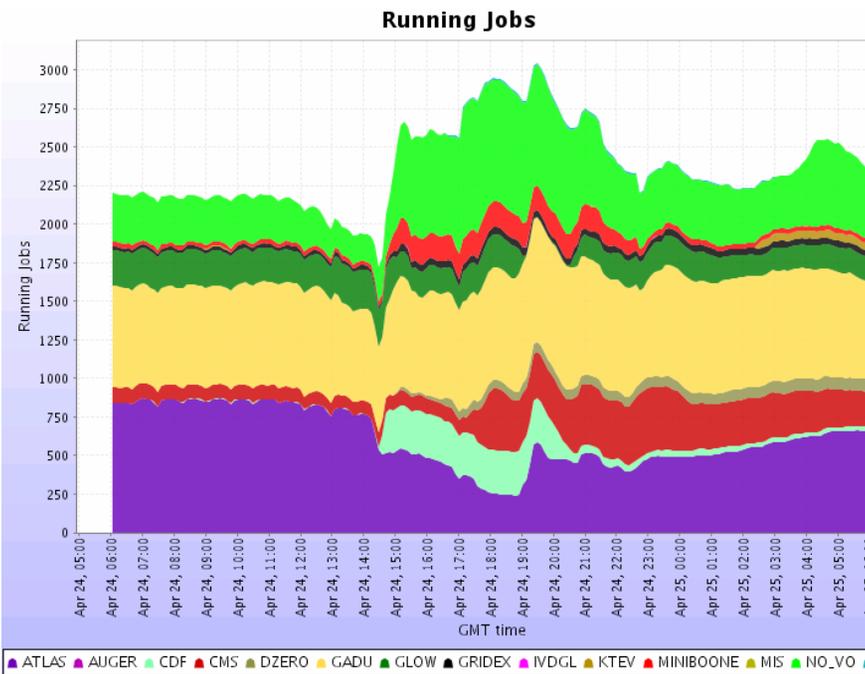
A Simple Analysis





First Generation Grids: On-Demand/Batch Computing

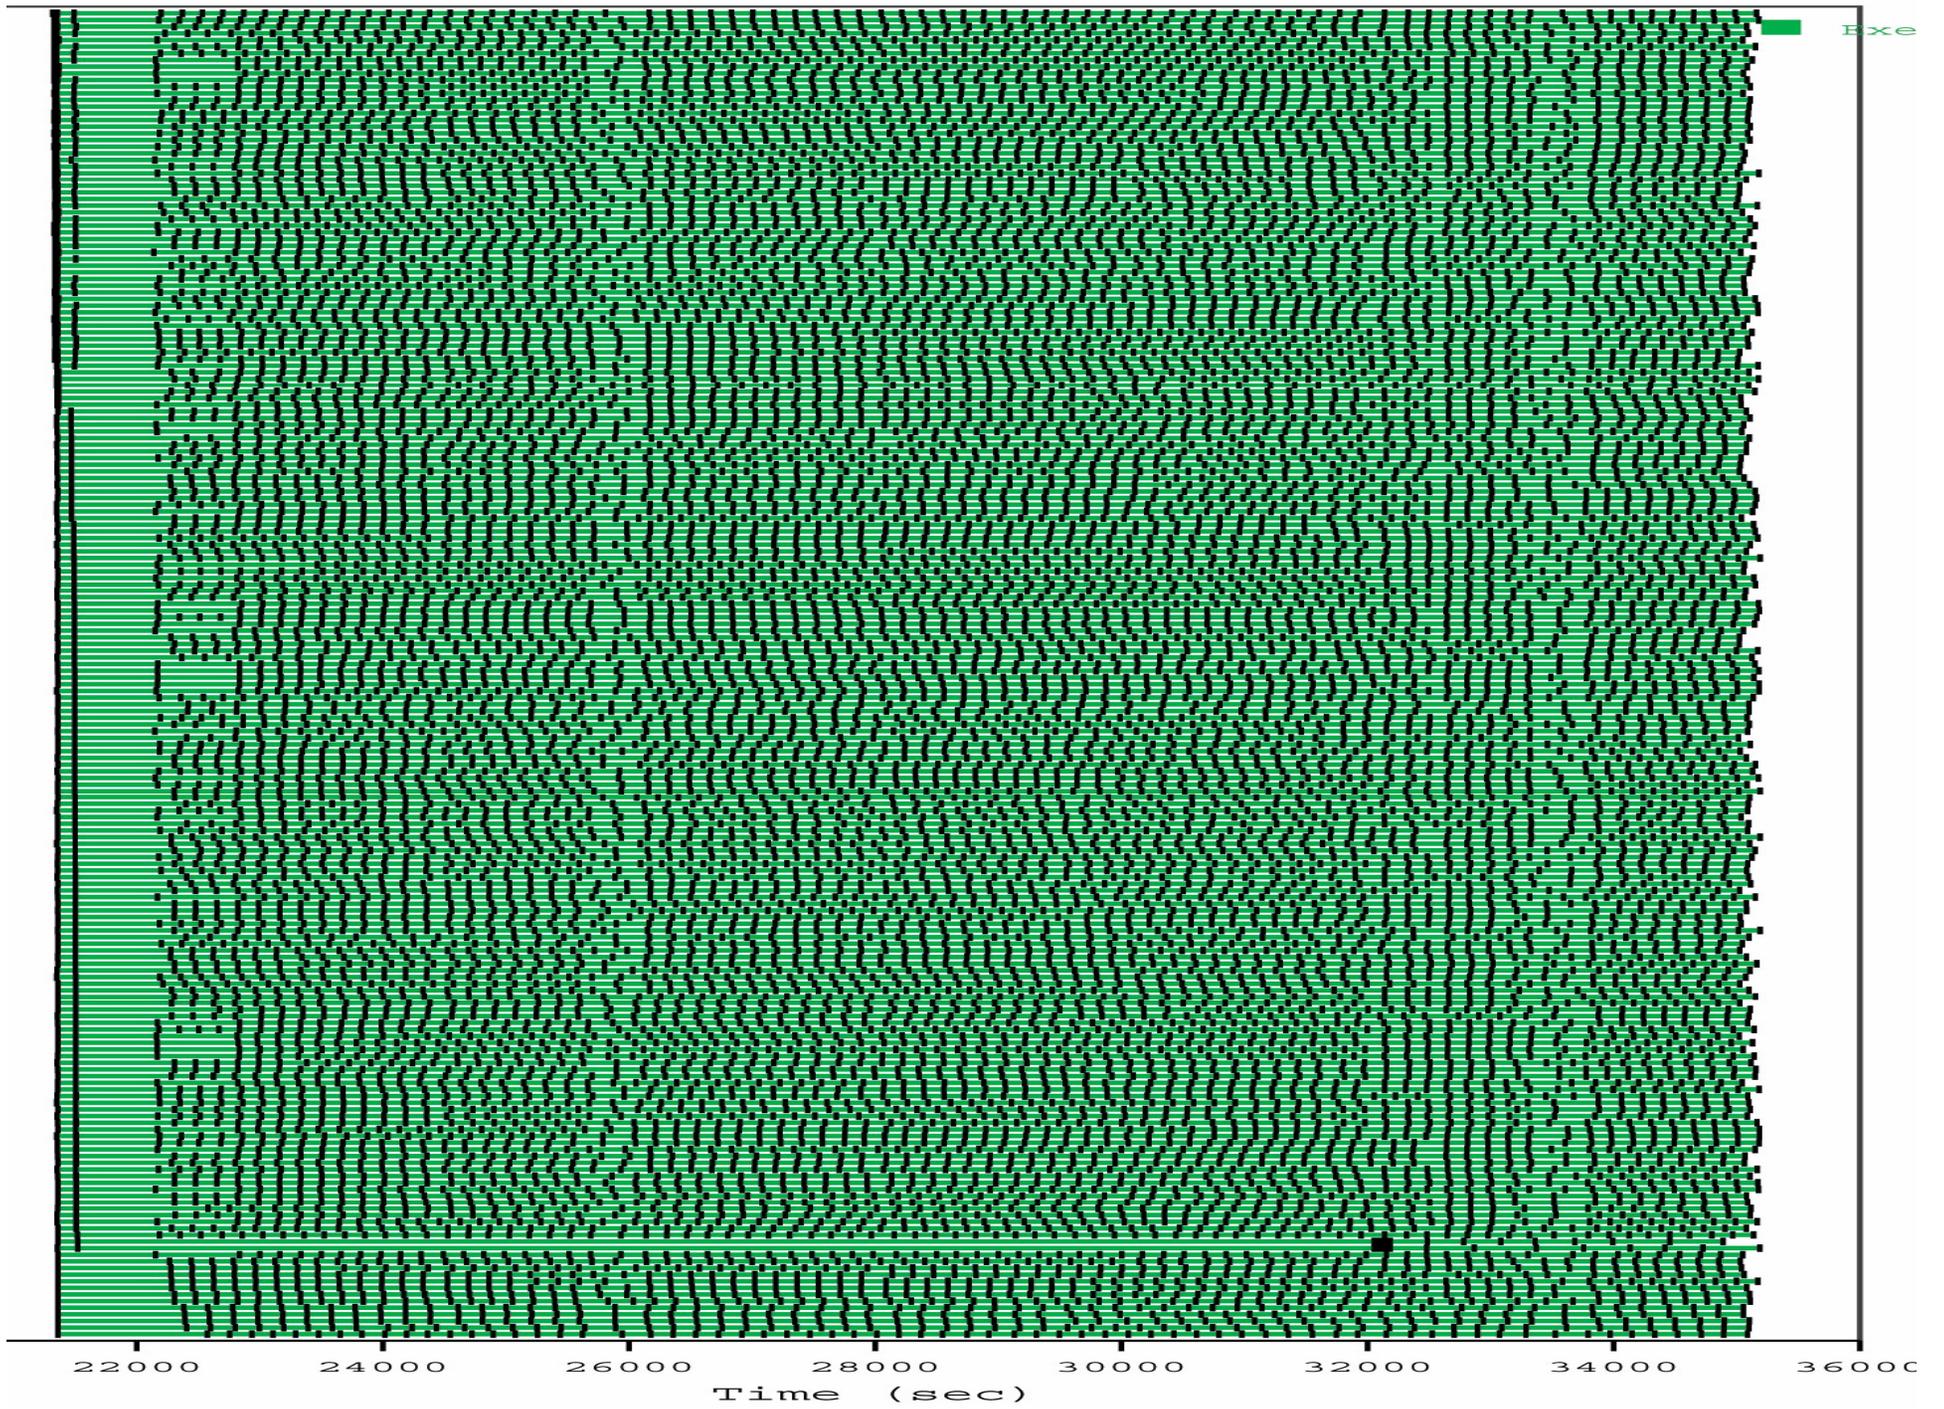
Focus on aggregation of many resources for
massively (data-)parallel applications

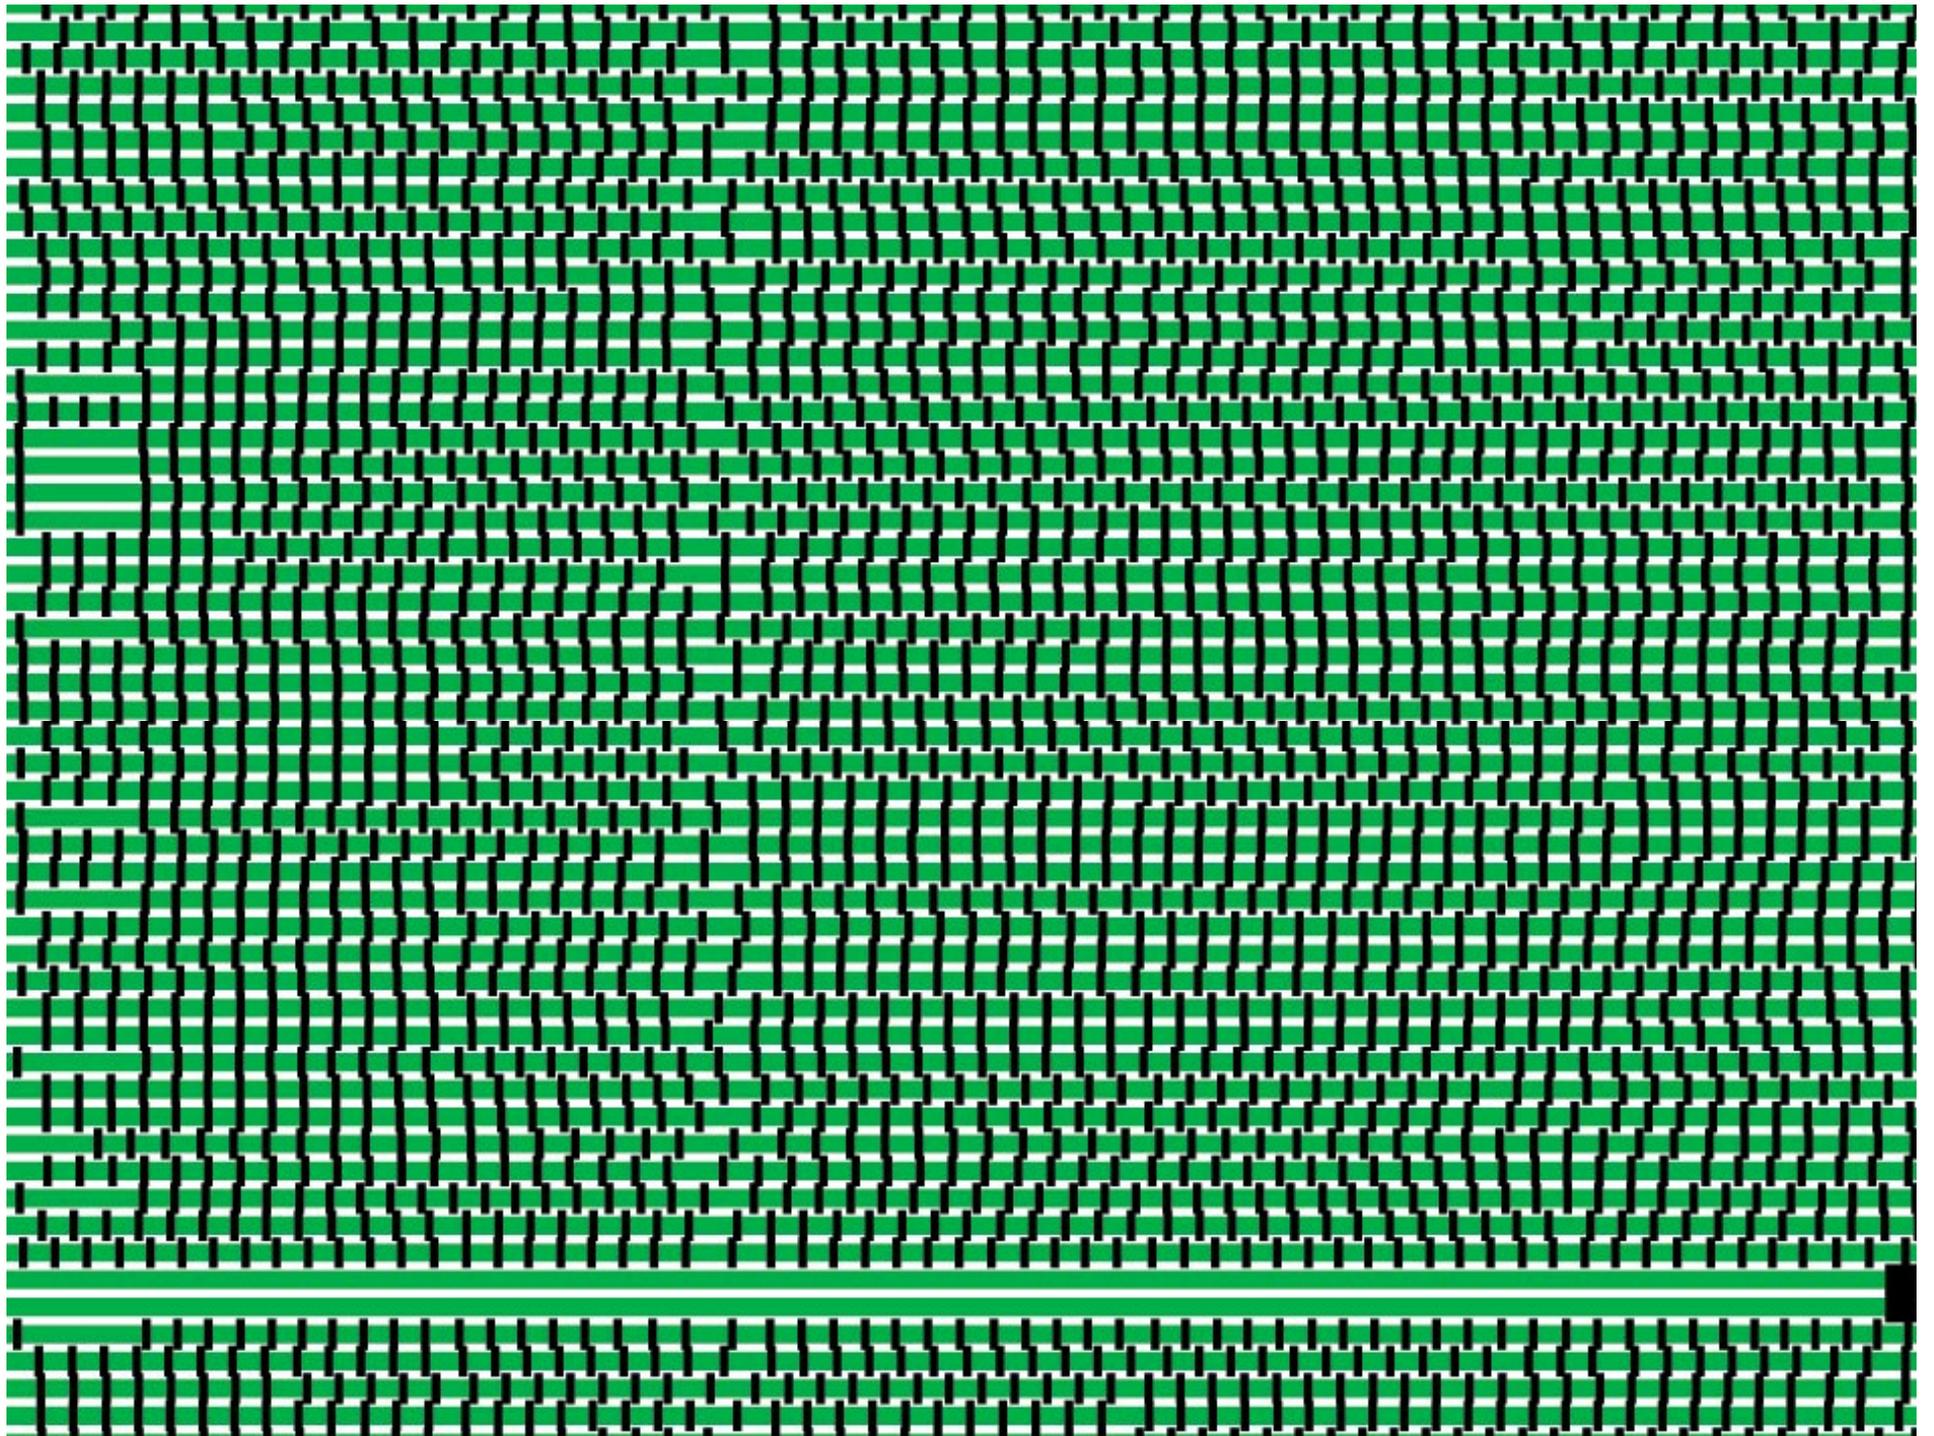


EGEE



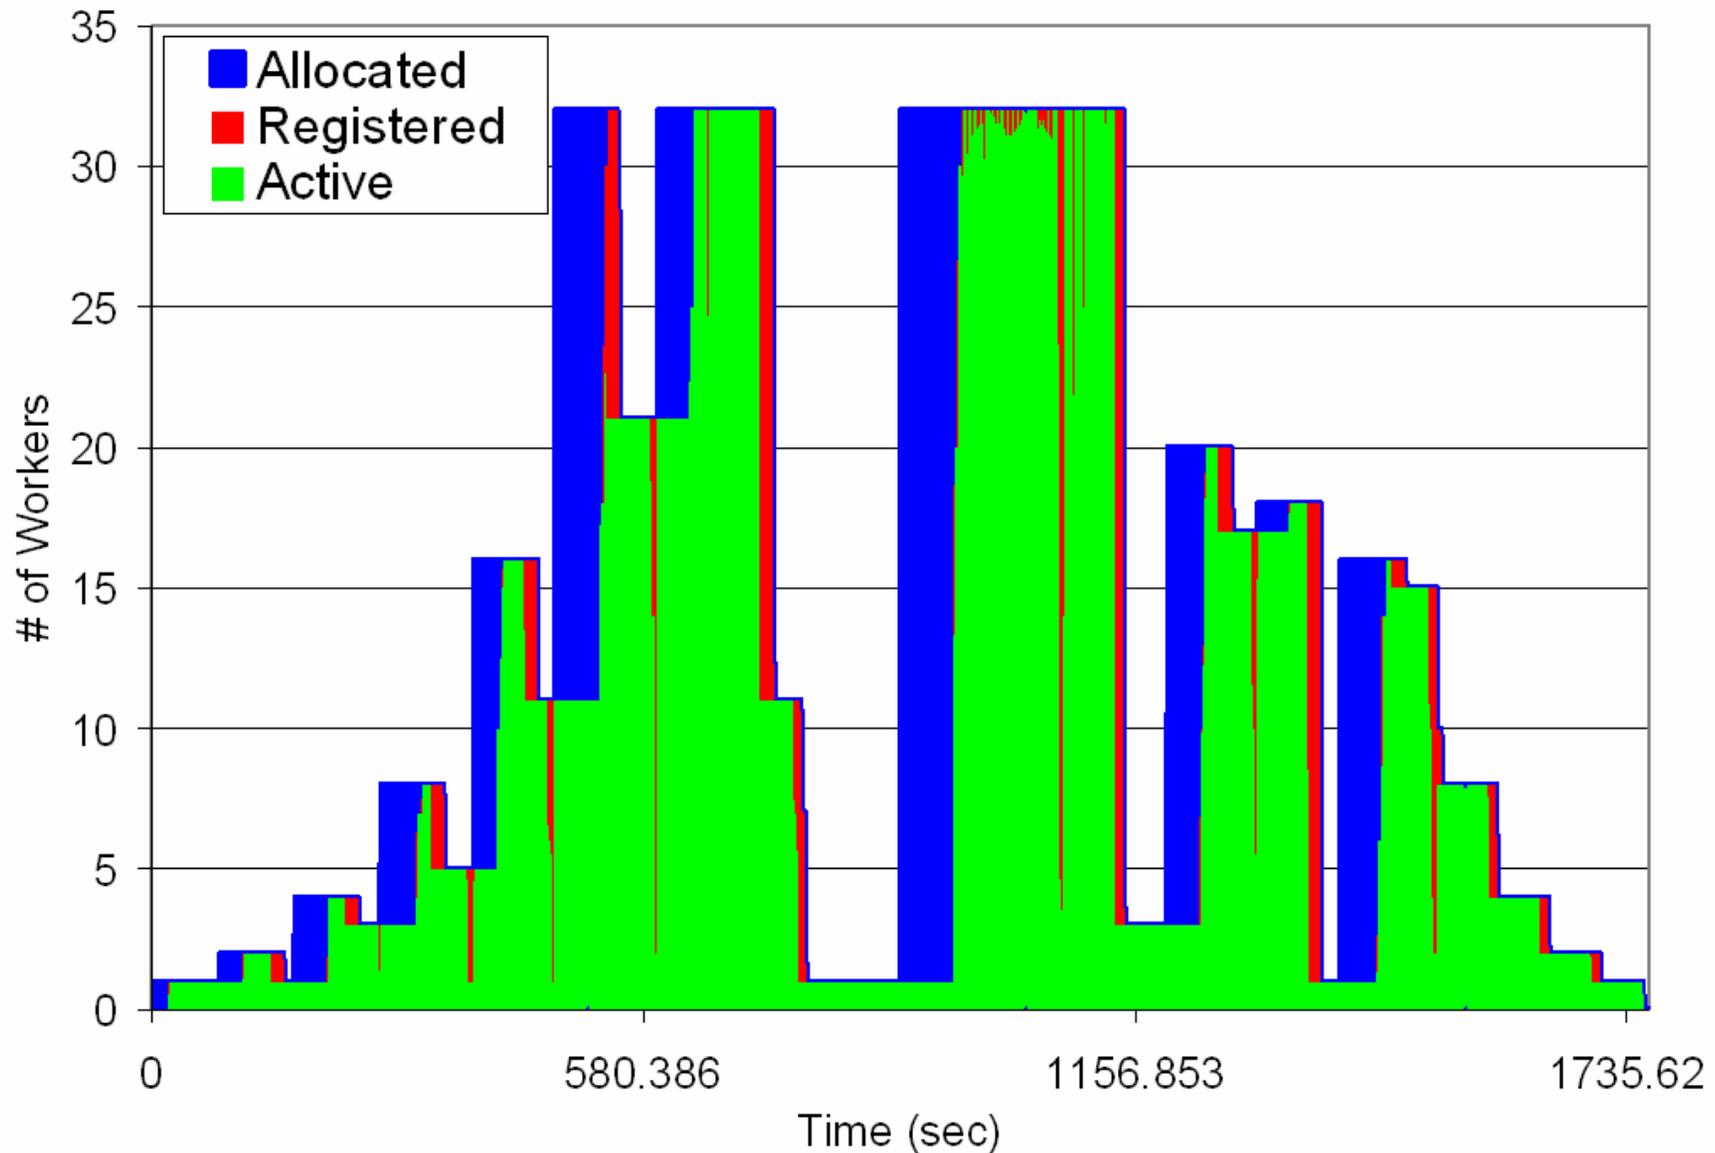
Globus 13





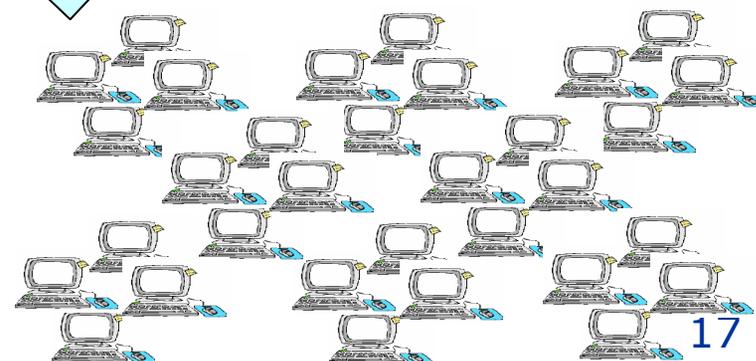
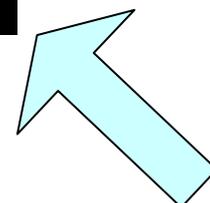
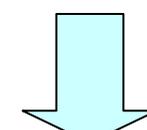
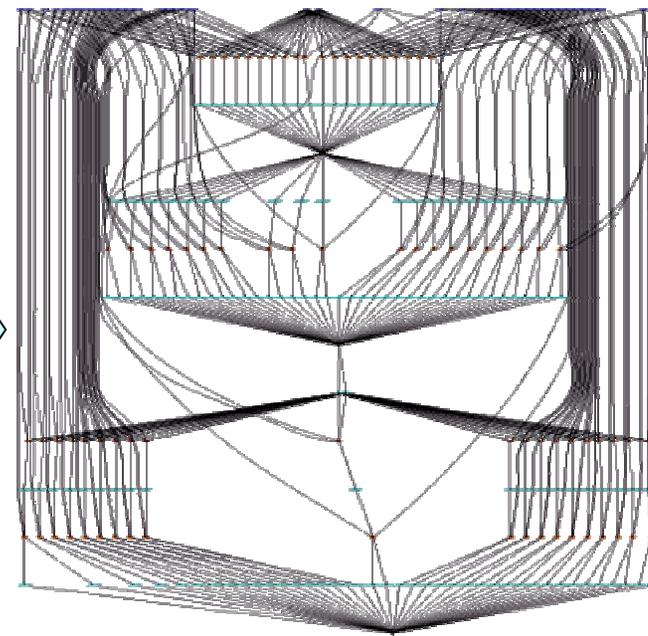
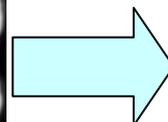
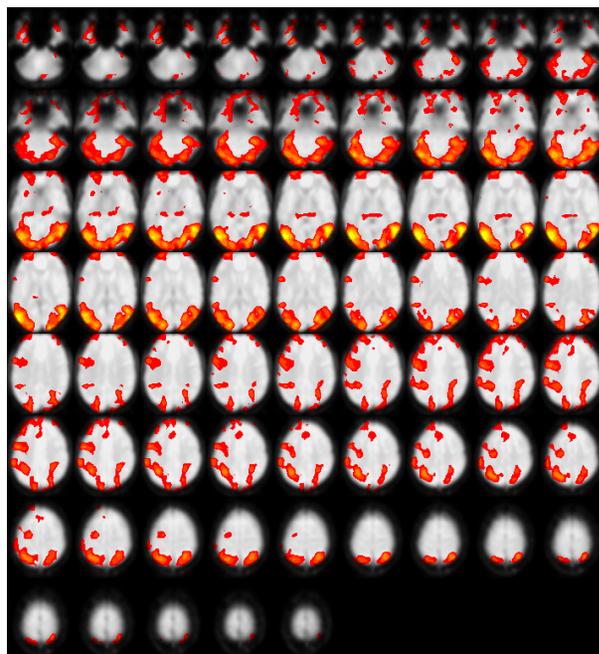


Dynamic Allocation for Dynamic Workloads





Many Users Analyze fMRI Data in Many Different Ways



- Wide range of analyses
 - ◆ Testing, interactive analysis, production runs
 - ◆ Data mining
 - ◆ Parameter studies



Second Generation Grids: **Service-Oriented Science**

- Empower many more users by enabling on-demand access to **services**
- Grids become an enabling technology for **service oriented science** (or business)
 - ◆ Grid infrastructures host services
 - ◆ Grid technologies used to build services



*Science
Gateways*

TeraGrid™
EMPOWERING DISCOVERY

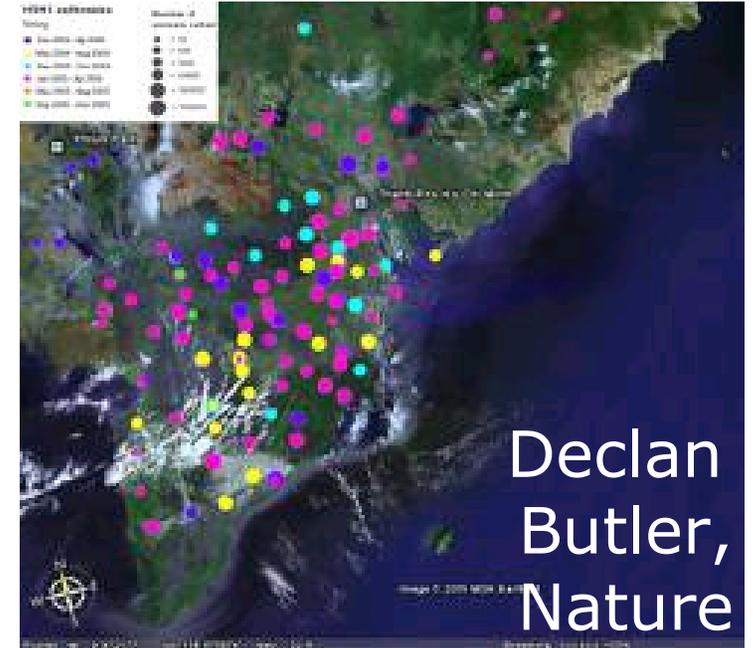


“Service-Oriented Science”, *Science*, 2005



“Web 2.0”

- Software as services
 - ◆ Data- & computation-rich network services
- Services as platforms
 - ◆ Easy composition of services to create new capabilities (“mashups”)—that themselves may be made accessible as new services
- Enabled by massive infrastructure buildout
 - ◆ Google projected to spend \$1.5B on computers, networks, and real estate in 2006
 - ◆ Many others are spending substantially
- Paid for by advertising





Service-Oriented Science: E.g., Virtual Observatories

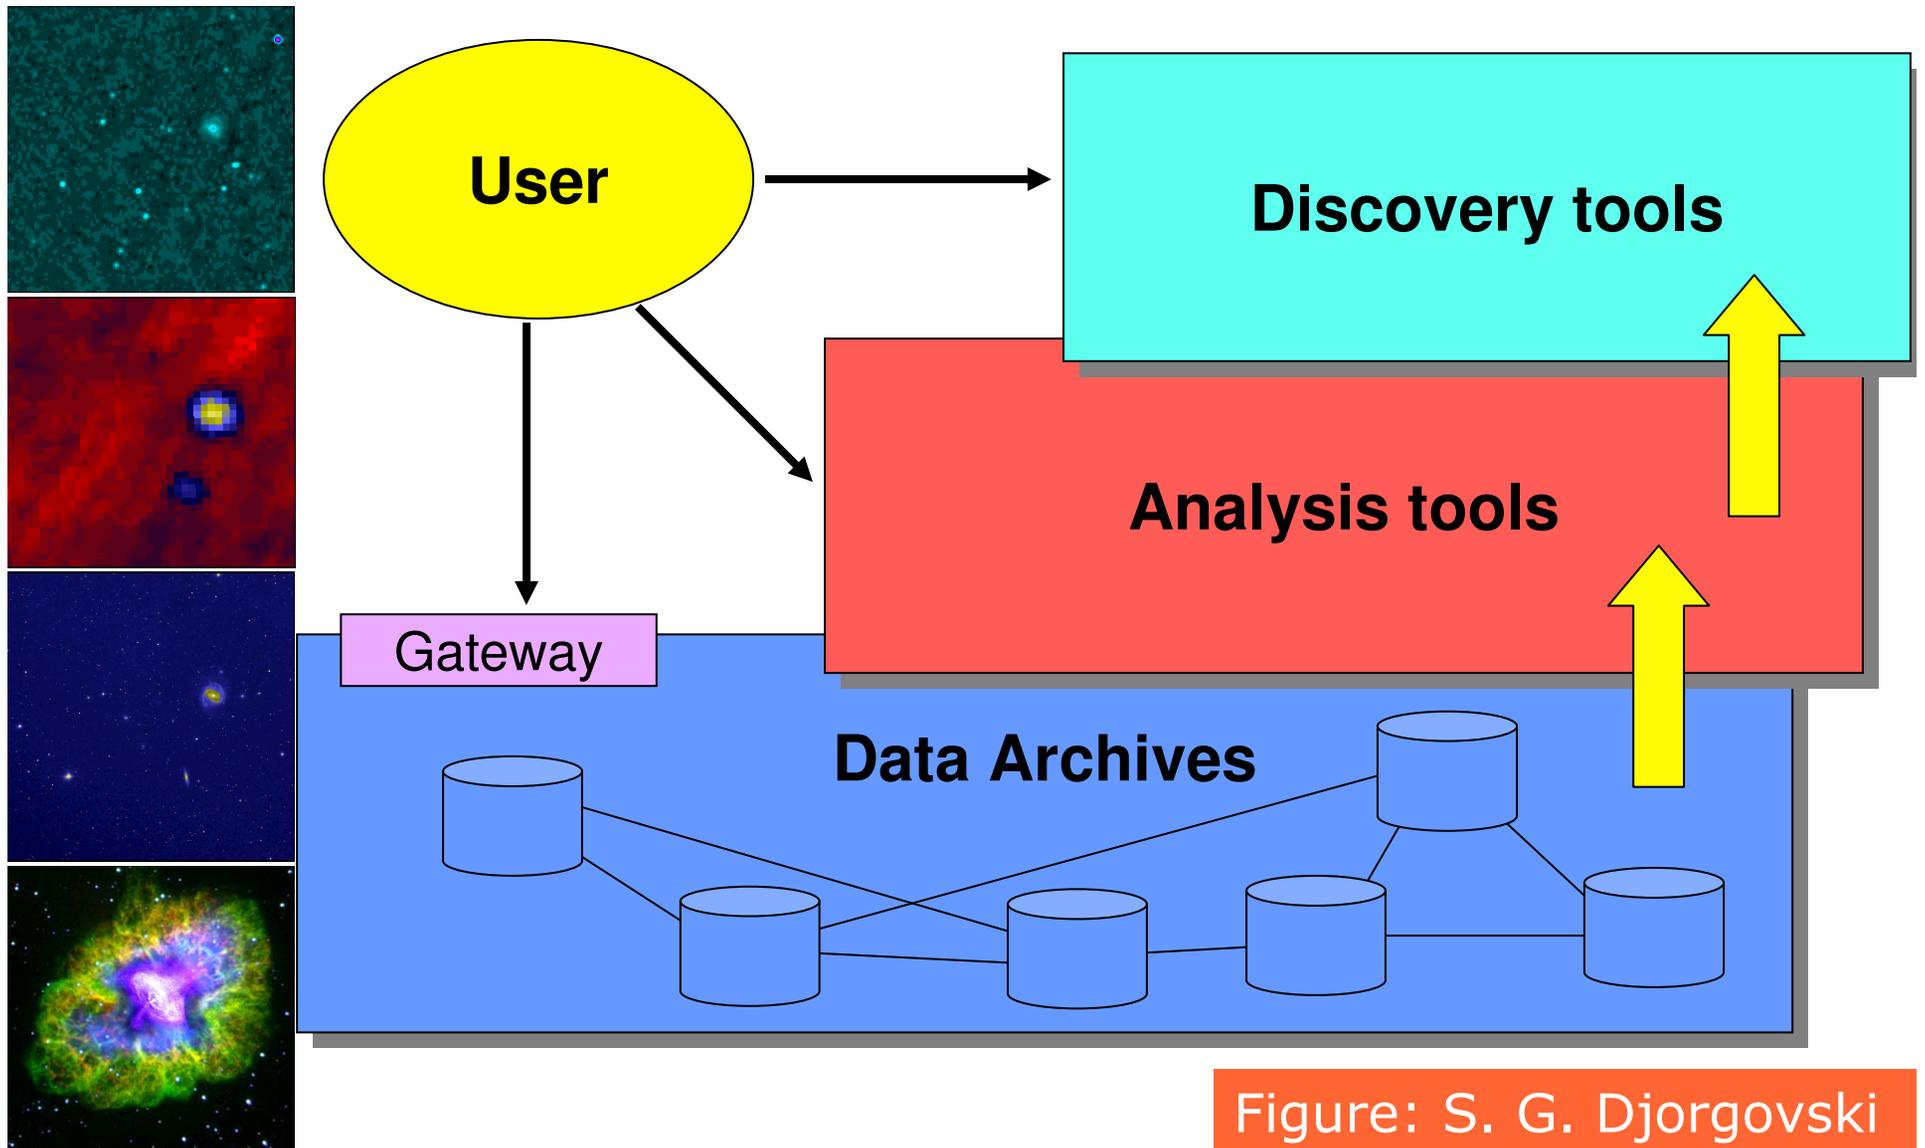
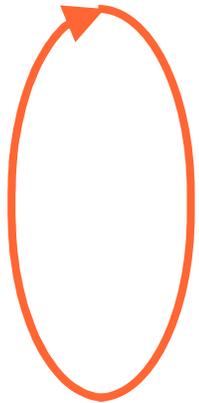


Figure: S. G. Djorgovski



Service-Oriented Science



People **create** services (data or functions) ...
which I **discover** (& decide whether to use) ...
& **compose** to create a new function ...
& then **publish** as a new service.

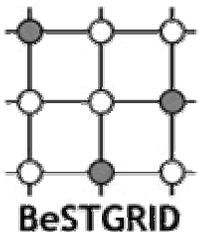
→ I find "someone else" to **host** services,
so I don't have to become an expert in
operating services & computers!



TeraGrid™
EMPOWERING DISCOVERY



→ I hope that this "someone else" can
manage security, reliability, scalability, ...



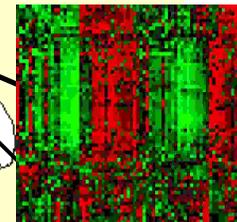
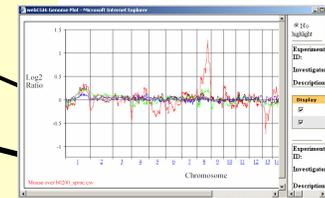
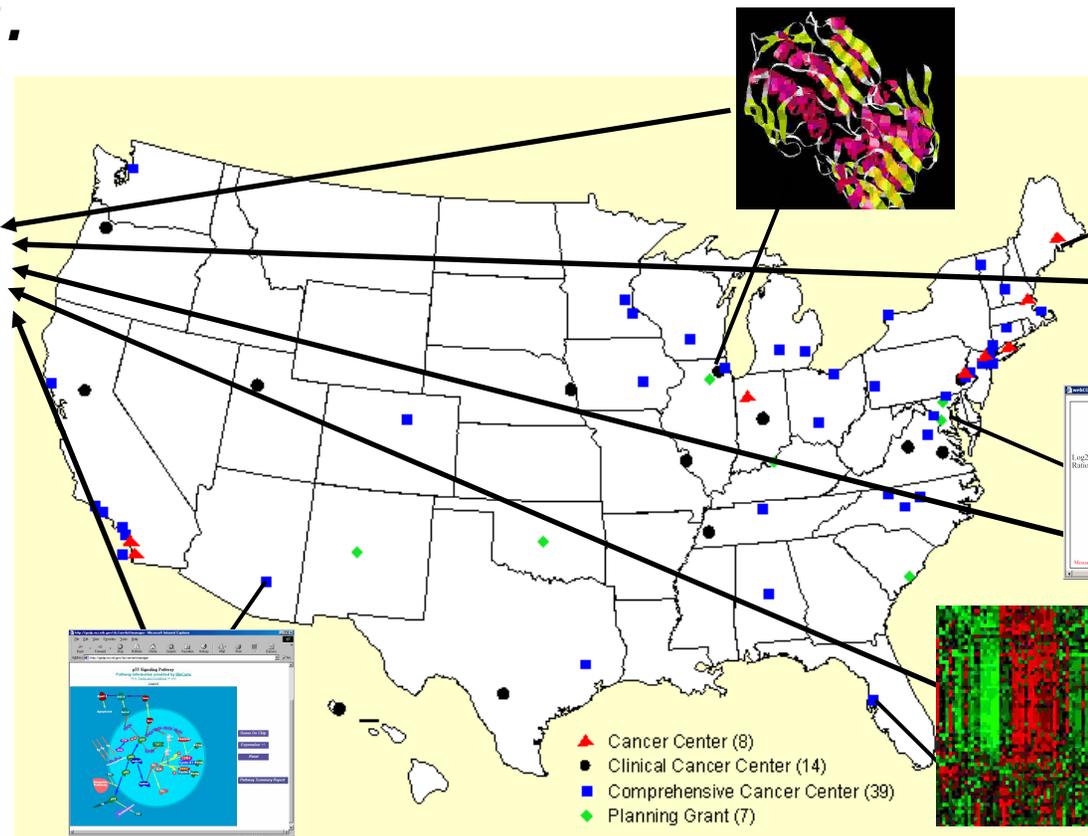
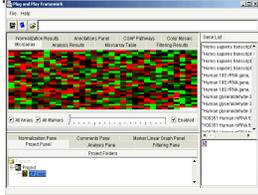
BeSTGRID

"Service-Oriented Science", *Science*, 2005

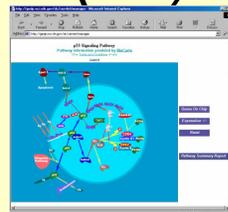


Service-Oriented Science & Cancer Biology

*caBIG: sharing of infrastructure, applications,
and data.*



**Data
Integration!**



- ▲ Cancer Center (8)
- Clinical Cancer Center (14)
- Comprehensive Cancer Center (39)
- ◆ Planning Grant (7)



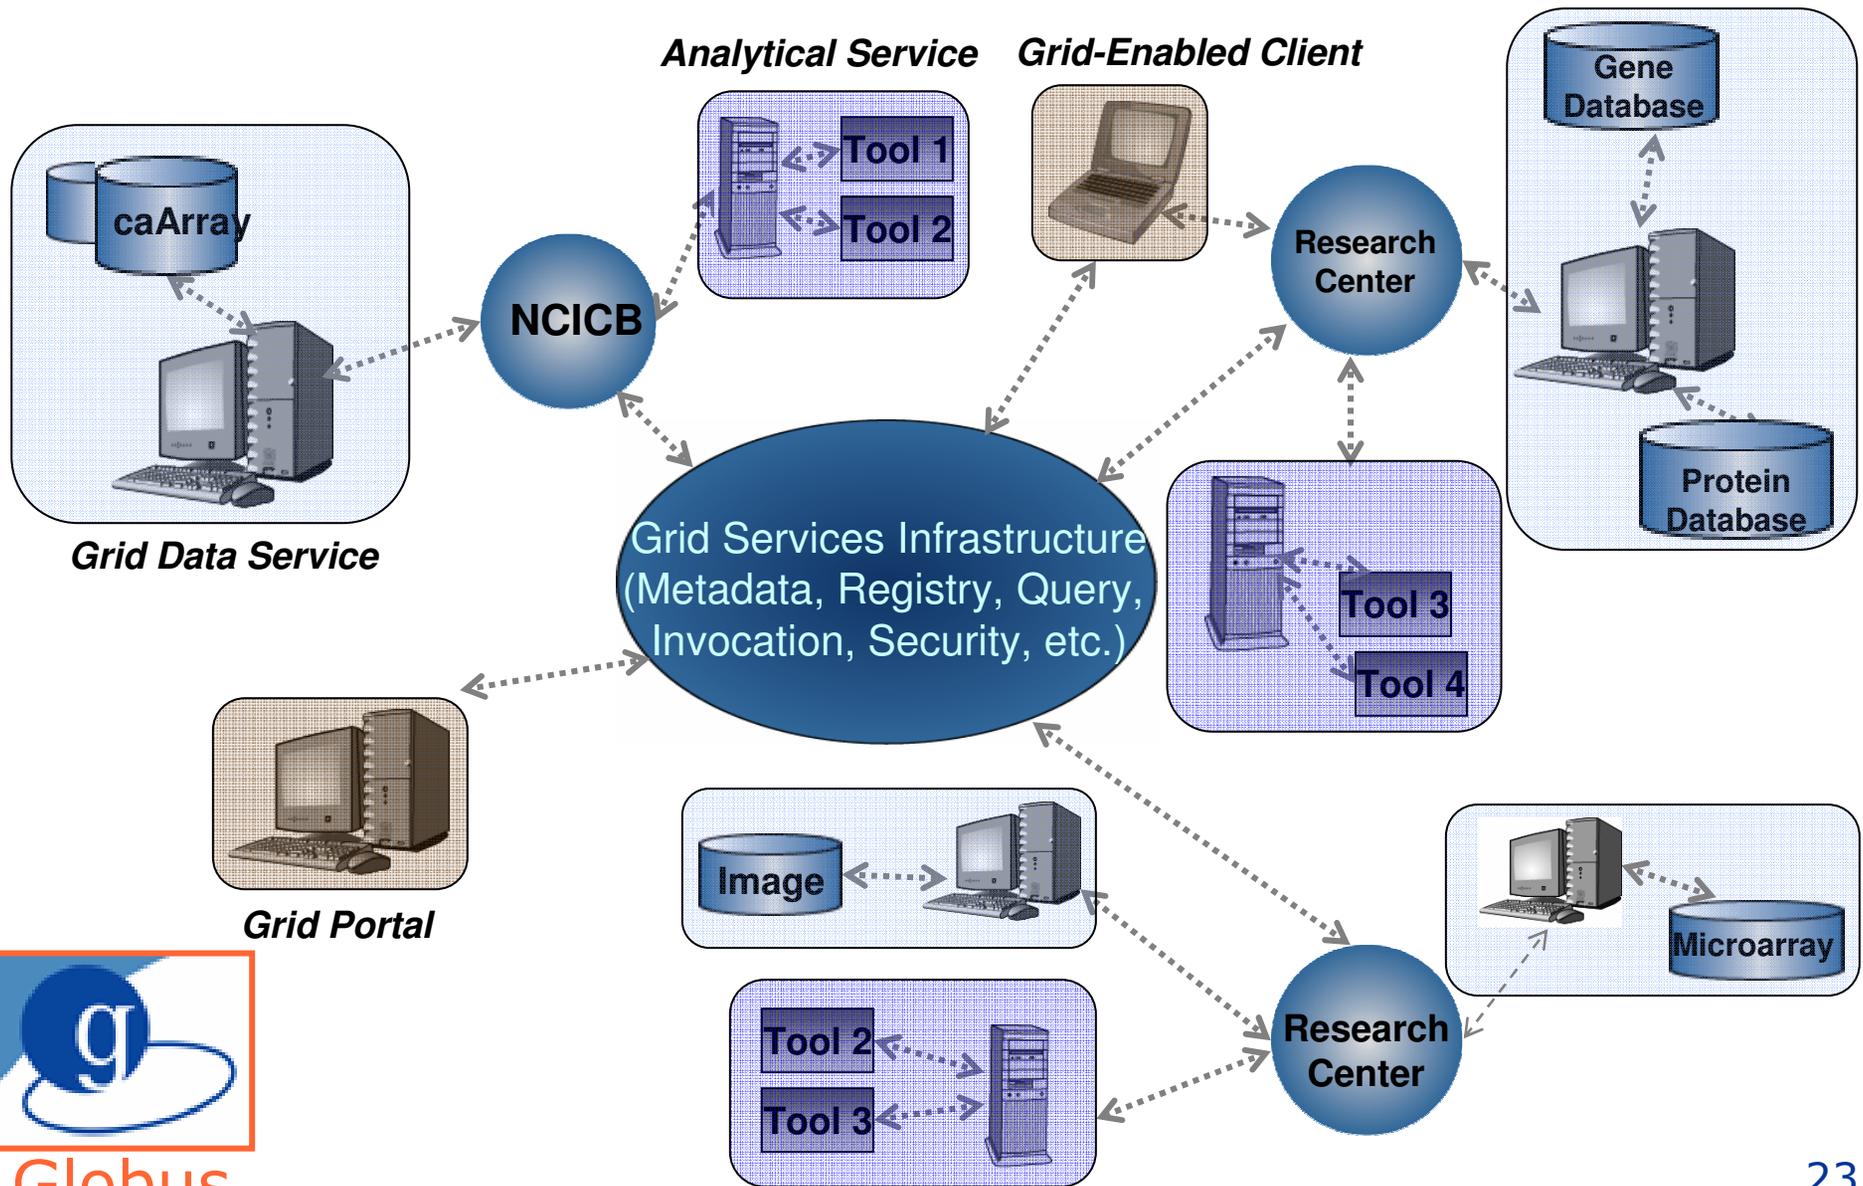
caBIG

cancer Biomedical
Informatics Grid





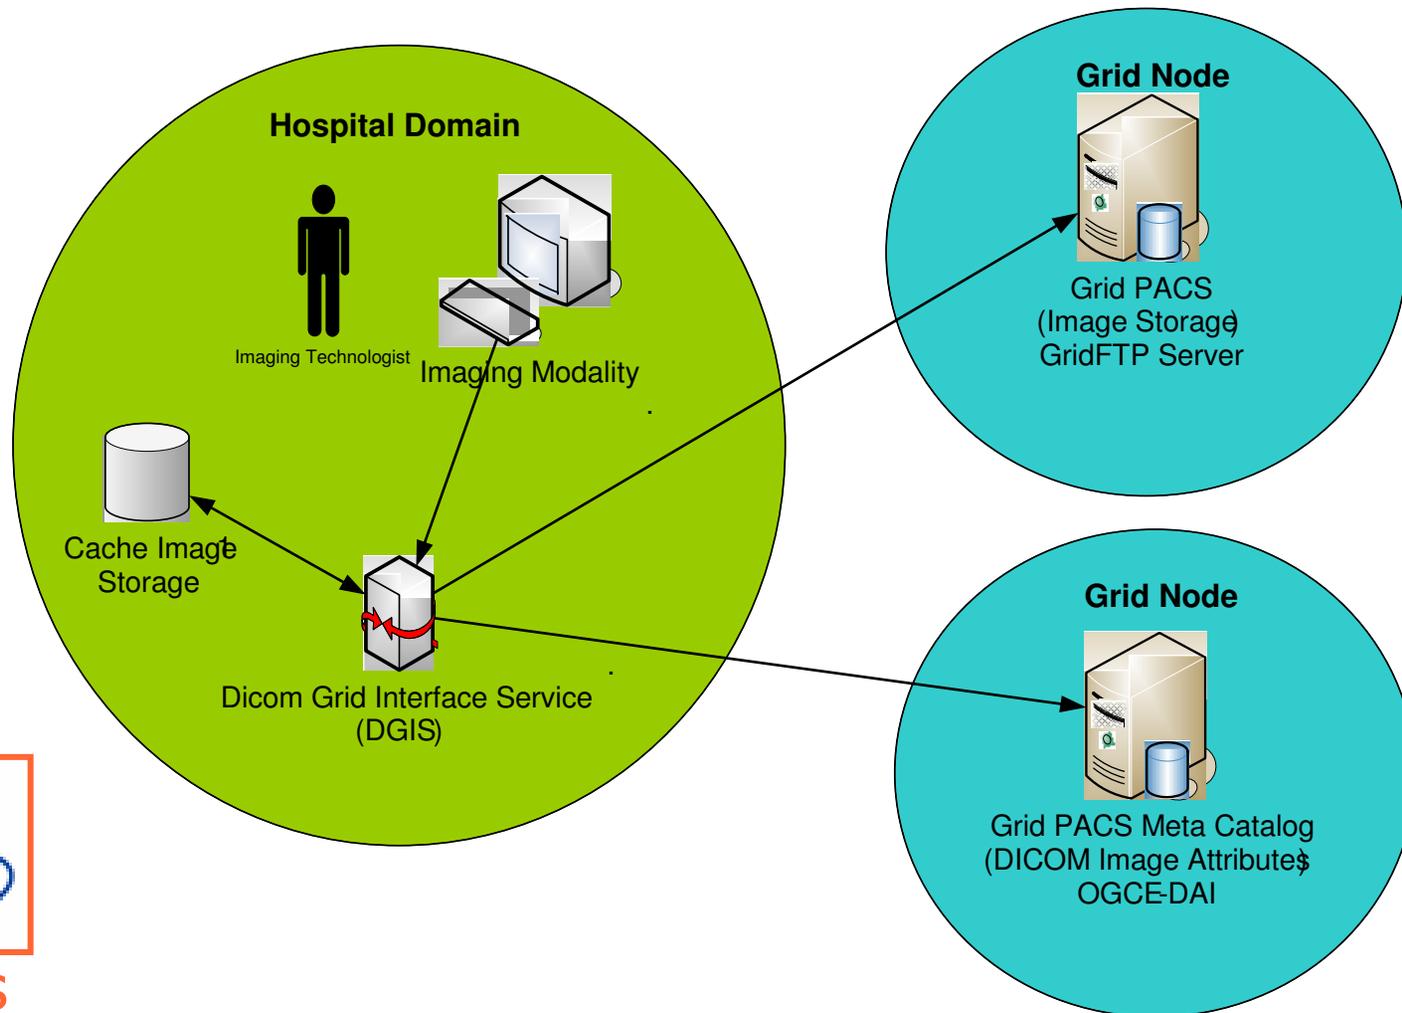
Cancer Bioinformatics Grid



Globus



MEDICUS: Management of DICOM Images

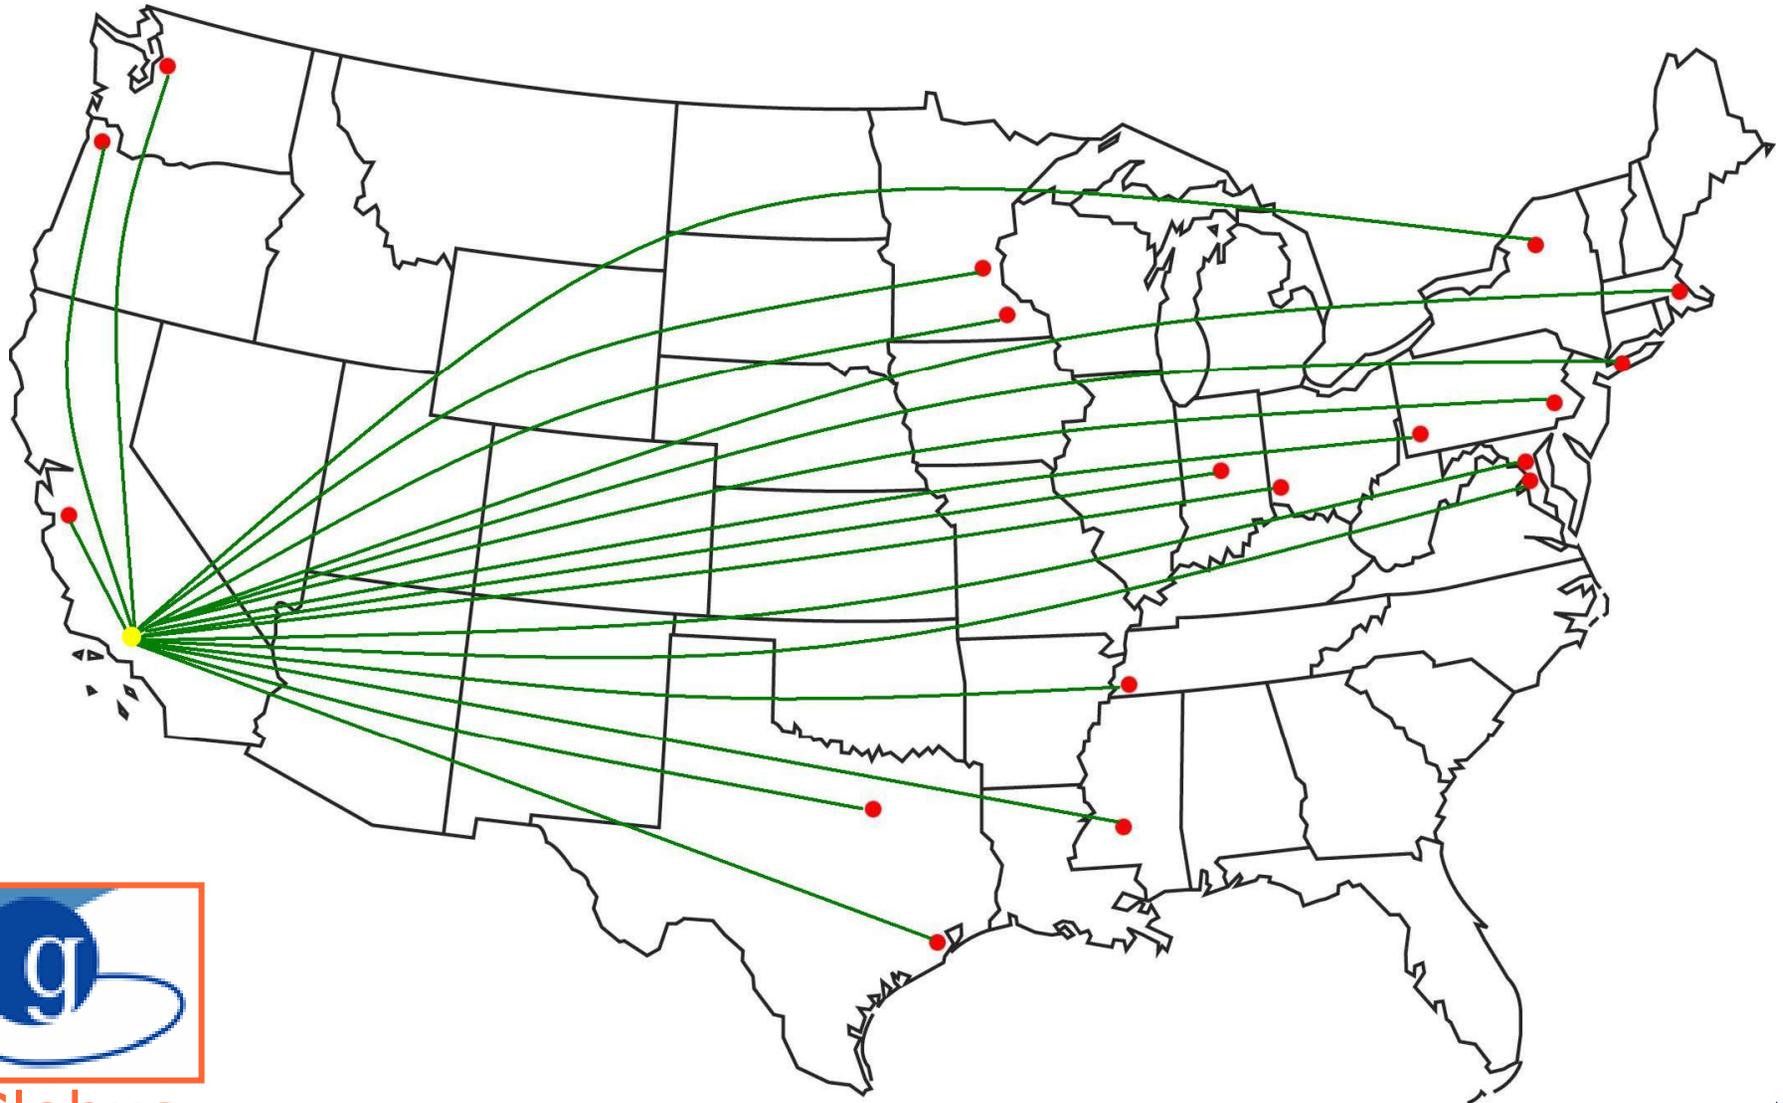


Globus

Stephan Erberich, Manasee Bhandekar, Ann Chervenak, et al.



Children's Oncology Grid: A MEDICUS Deployment



Globus



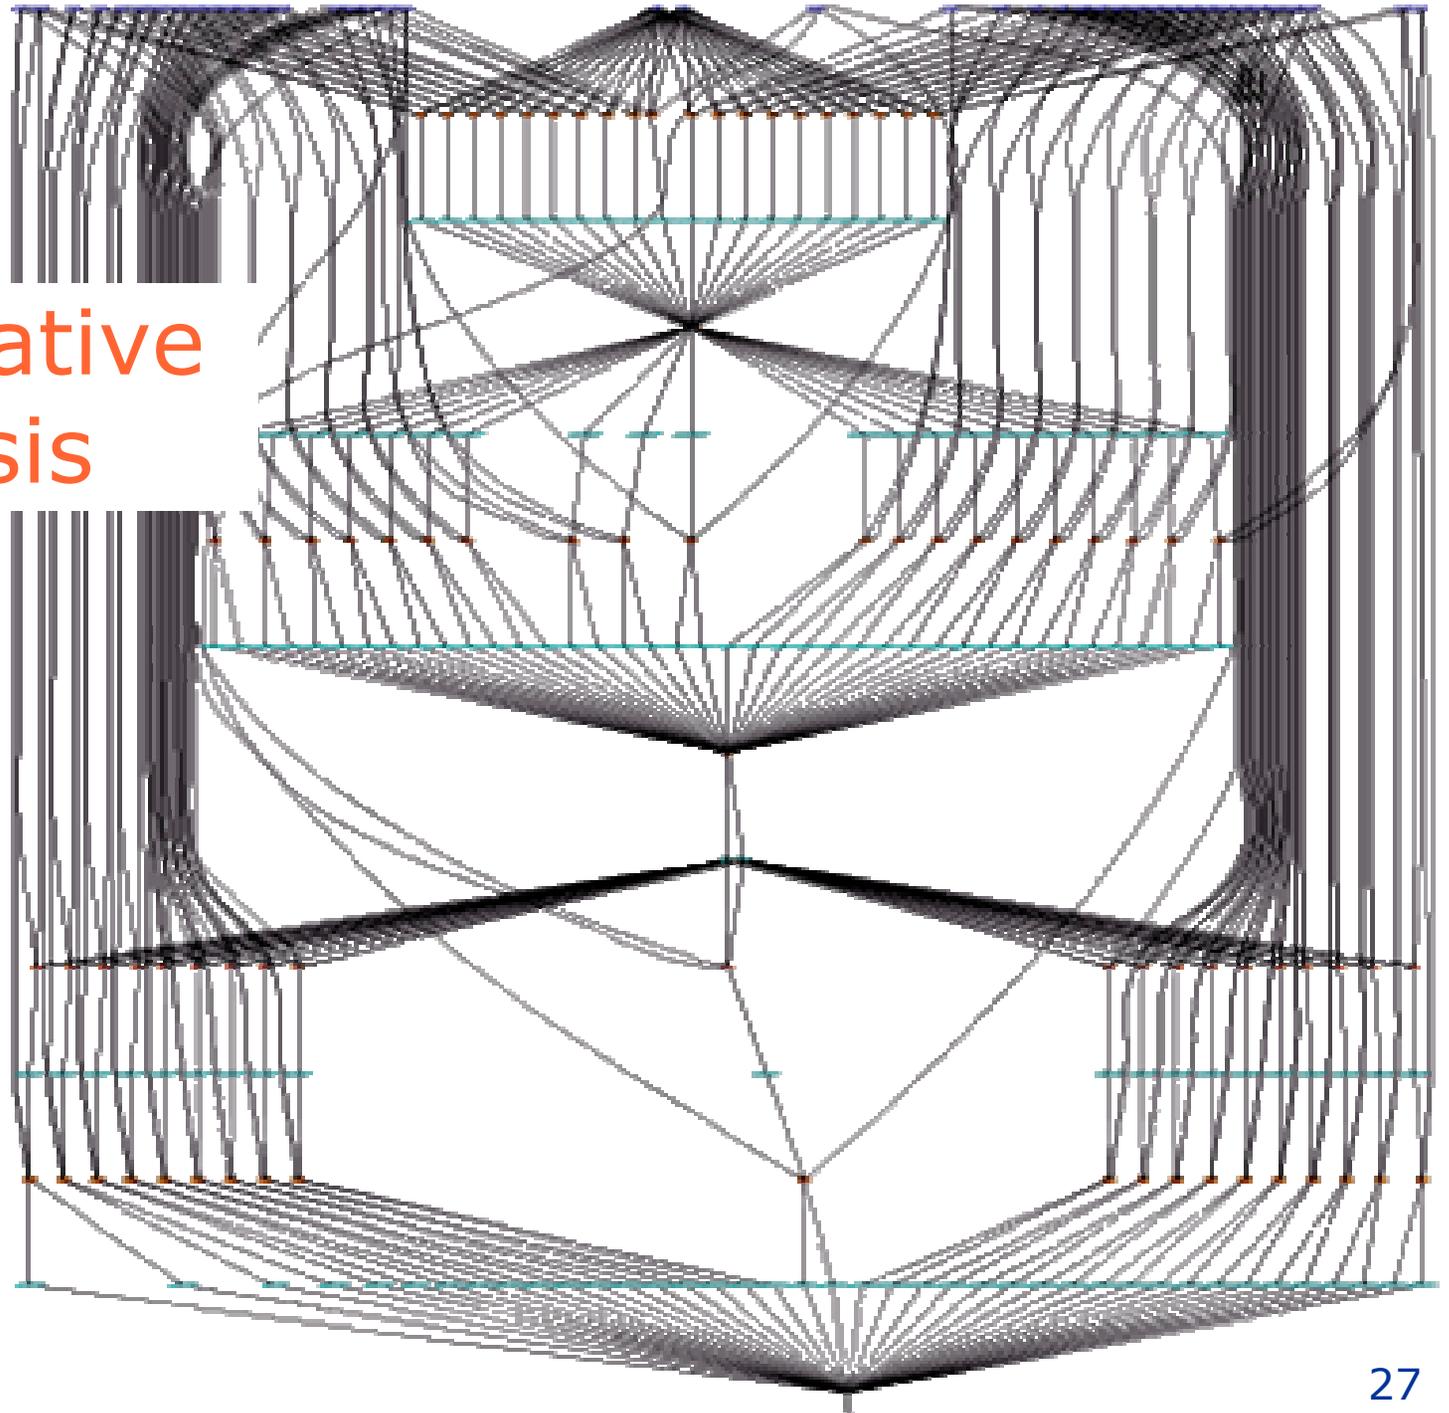
MEDICUS Under the Covers

Globus Toolkit Release 4

- DICOM images
 - ◆ Send (publish)
 - ◆ Query/Retrieve (discover)
 - Grid Archive
 - ◆ Fault tolerant
 - ◆ Bandwidth
 - Security
 - ◆ Authentication
 - ◆ Authorization
 - ◆ Cryptography
 - Access
 - ◆ Web portal
 - Applications
 - ◆ Computing
 - ◆ Data Mining
- *DICOM Grid Interface Service (DGIS)*
+
Meta Catalog Service (OGSA-DAI)
- *Data Replication Service (DRS)*
- *X.509 Certificates*
+
MyProxy Delegation
- *Grid Web Portal, OGCE / GridSphere*
- *GRAM, OGSA-DAI*



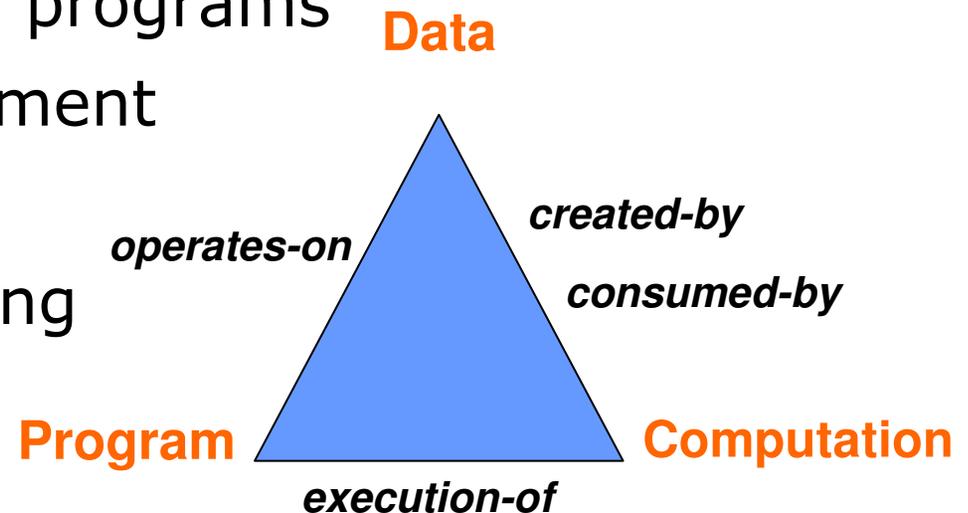
Collaborative Analysis





The Virtual Data Concept

- Capture information about relationships among
 - ◆ Data (varying locations and representations)
 - ◆ Programs (& inputs, outputs, constraints)
 - ◆ Computations (& execution environments)
- Apply this information to:
 - ◆ Discovery of data and programs
 - ◆ Computation management
 - ◆ Provenance
 - ◆ Planning and scheduling
 - ◆ Performance optimization





Provenance Model

- Temporal aspect

- ◆ **Prospective** provenance

- Recipes for how to produce data
 - Metadata annotations about procedures and data

- ◆ **Retrospective** provenance

- Invocation records of run time environments and resources used: site, host, executable, execution time, file stats ...

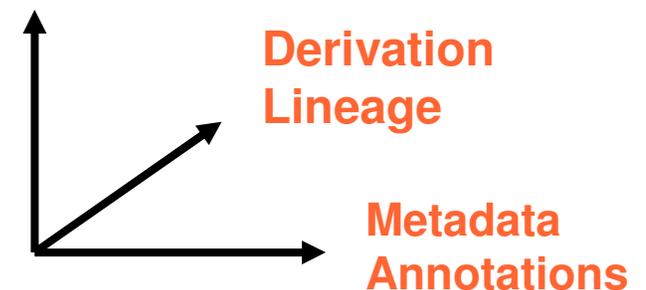
- Dimensional aspect

- ◆ **Virtual data** relationships

- ◆ **Derivation** lineage

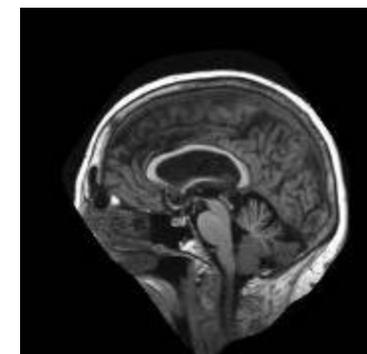
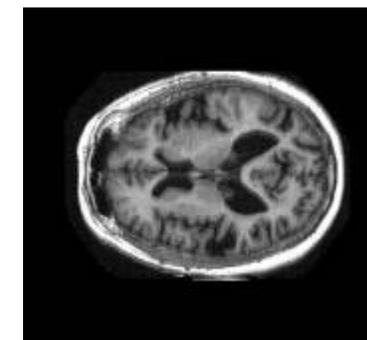
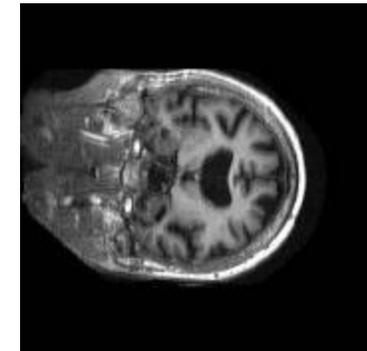
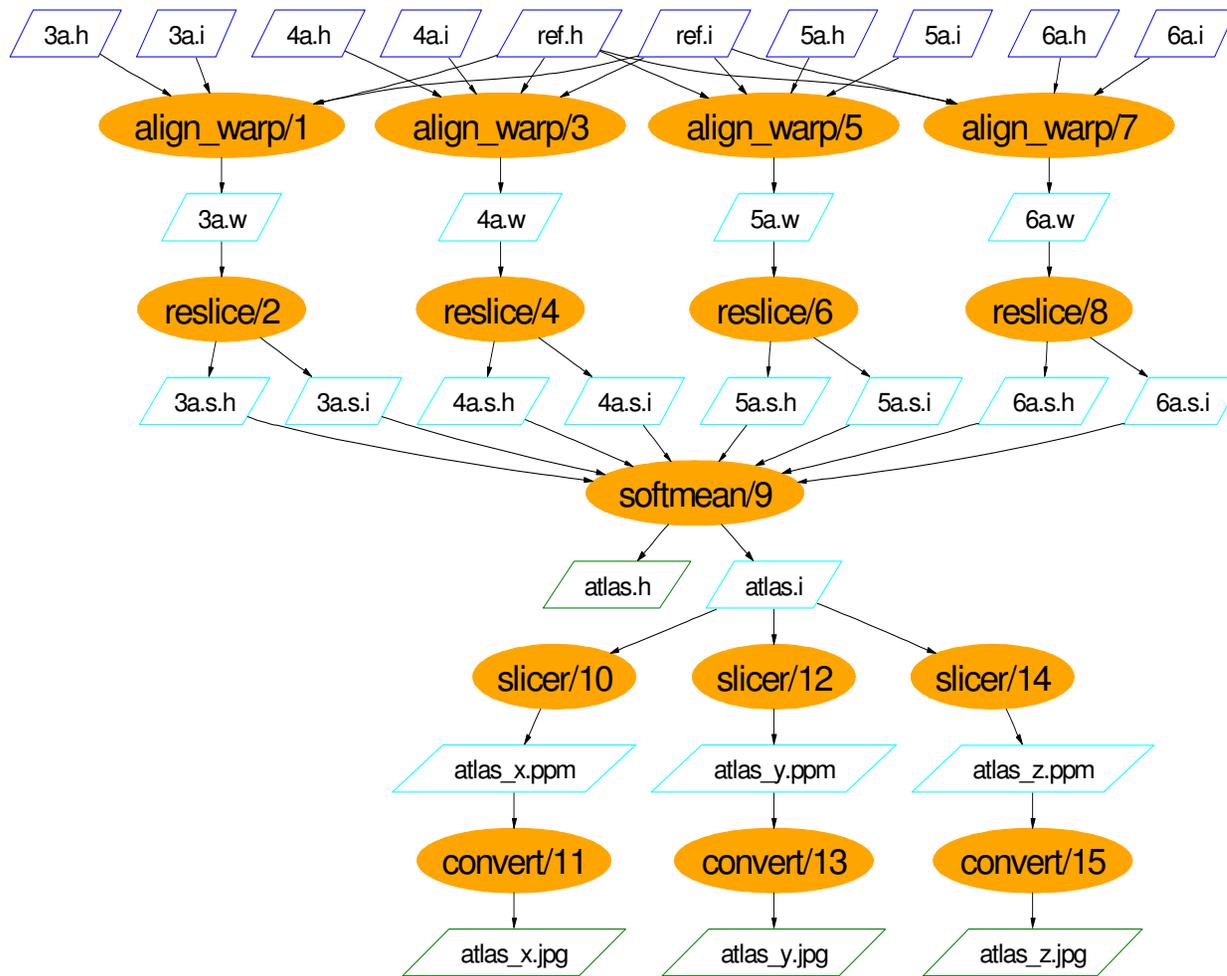
- ◆ **Metadata** annotations

Relationships





Query Context: fMRI Analysis



First Provenance Challenge, <http://twiki.ipaw.info/>



Query Examples

- Query by procedure signature
 - ◆ Show procedures that have inputs of type *subjectImage* and output types of *warp*
- Query by actual arguments
 - ◆ Show *align_warp* calls (including all arguments), with argument *model=rigid*
- Query by annotation
 - ◆ List anonymized subject images for young subjects:
 - Find datasets of type *subjectImage* , annotated with *privacy=anonymized* and *subjectType=young*
- Basic lineage graph queries
 - ◆ Find all datasets derived from dataset '5a'
- Graph pattern matching
 - ◆ Show me all output datasets of *softmean* calls that were aligned with *model=affine*
- Multi-dimensional query

Thailand joins the grid

By [Don Sambandaraksa](#), Bangkok Post

Wednesday, May 16 2007 11:34 AM



Bangkok Post

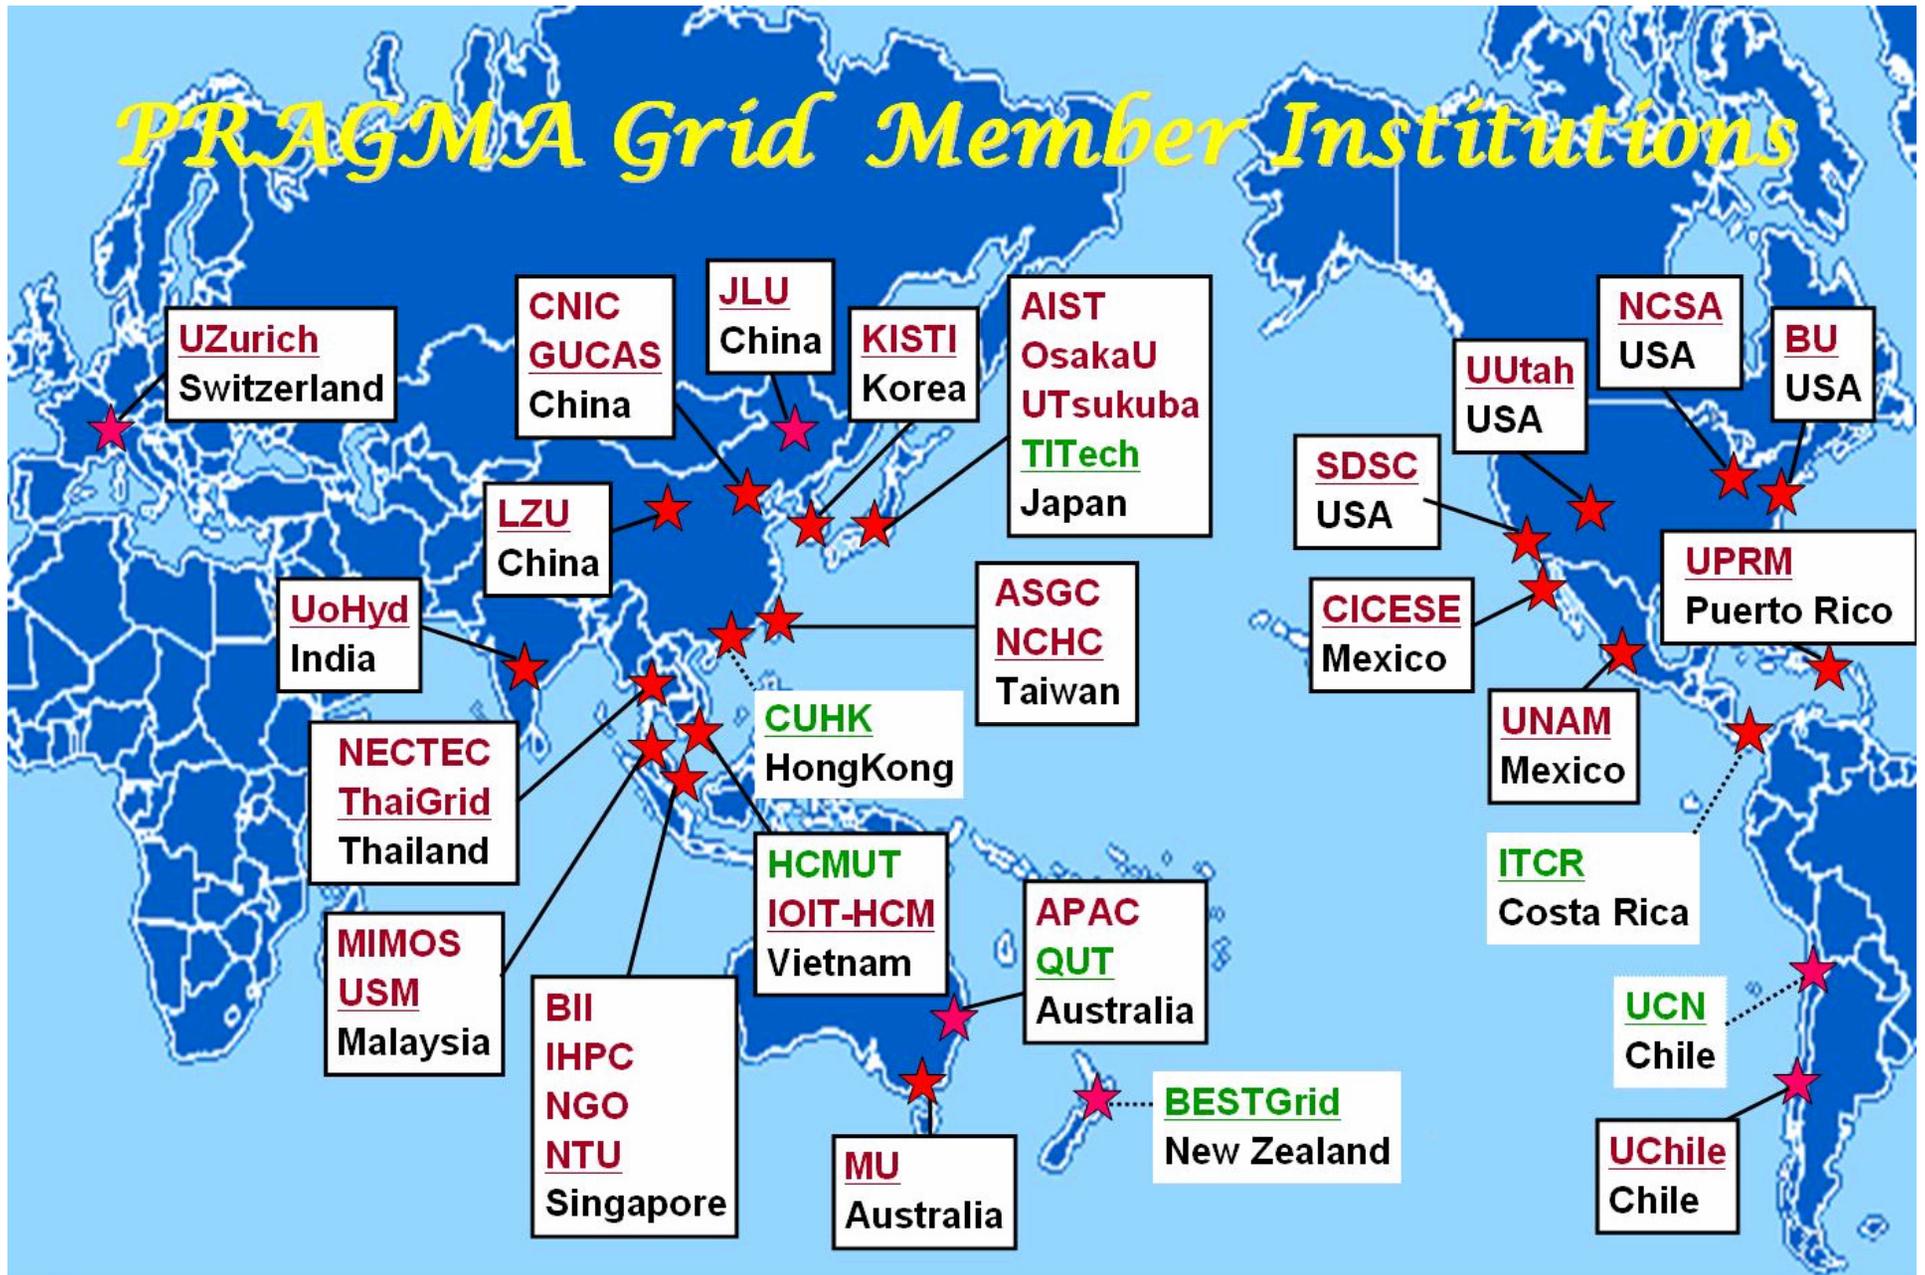
www.bangkokpost.com

Grid computing is no longer just about universities getting together to pool and share their computing power, but increasingly it is about grids of sensors, information and of experts that are shared and leveraged across the world. And now Thailand is playing a major part in this new wave, having recently hosted the 12th meeting of the Pacific Rim Application and Middleware Assembly (Pragma).

“Grids are not just communities of computers, but communities of researchers, of people.”

— Peter Arzberger

PRAGMA Grid Member Institutions



31 institutions in 15 countries/regions (+ 7 in preparation)

Last update: 5/30/2007





Globus Downloads Last Month





Grids are Communities of **People** as well as Computers

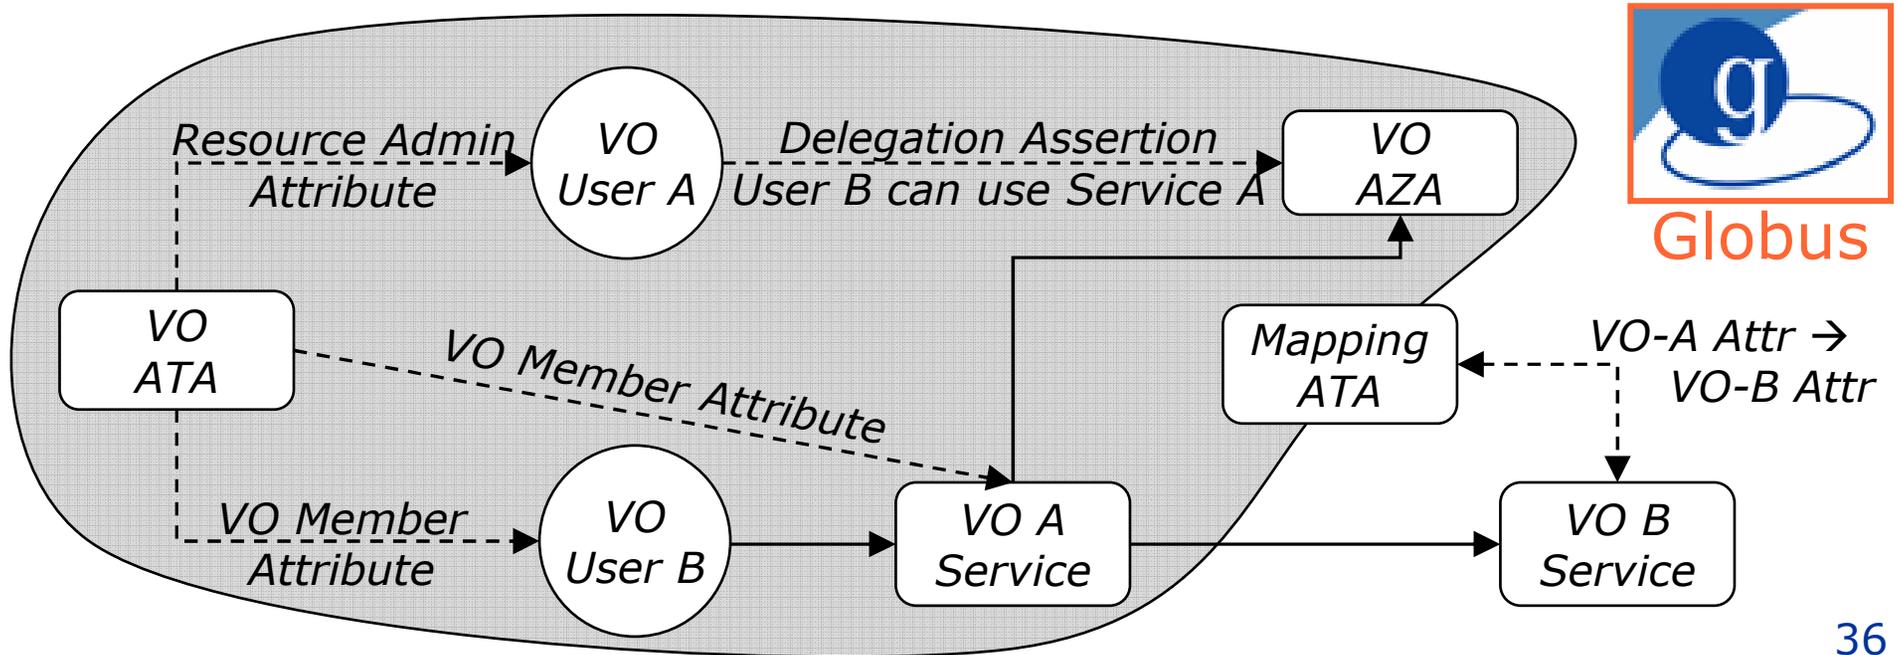
- Based on (technology-mediated) trust
 - ◆ Common goals
 - ◆ Processes and policies
 - ◆ Reward systems
- That share resources
 - ◆ Computers
 - ◆ Data
 - ◆ Sensor networks
 - ◆ Services
- Supported by software, standards, infrastructure





Security Services for Virtual Organization Policy

- Attribute Authority (ATA)
 - ◆ Issue signed attribute assertions (incl. identity, delegation & mapping)
- Authorization Authority (AZA)
 - ◆ Decisions based on assertions & policy
- Use with message- or transport-level security



Global Observation Database (View) Graph Observer

Social Informatics Data Grid

CR-Style Control Panel

Animated Text Transcript (Paragraph Representation)

Tag Transcript Editor

Animated Avatar Representation

Animated Graph Panes

Video Displays

Video List

Bennett Berthenthal et al., www.sidgrid.org

Five reasons to use Connotea

- [Save and organize](#) links to your references
- Easily [share](#) references with colleagues
- Access references from [any computer](#)
- [One click](#) is all it takes
- [Easy to use](#). Start creating your library today

[Learn more](#)

[Watch a short video](#) (2m 41s)

**Free online reference management
for all researchers, clinicians and
scientists**

Completely free, no download

[sign up](#)

[nature.com](#)

How it works

1 You find a reference

[Molecular structure of nucleic acids: a](#)
W JD, FH CRICK - Nature, 1953 - [View on Nature](#). 1953 Apr 25;171(4356):737-8 ...
[Web Search](#)

2 One click saves it on Connotea

Comments [Add To Connotea](#)

ature of nucleic acids: a

3 Add your own keywords for easy retrieval

DNA, genetics, geno

[genomics](#)

Explore

[Home](#)
[Recent Activity](#)
[Popular Links](#)
[Popular Tags](#)

Beginner's guide

[About Connotea](#)
[How It Works](#)
[Getting Started](#)
[Introductory Videos](#)

Account details

[My Account](#)
[Advanced Settings](#)
[My Library](#)
[My Community Profile](#)

Help

[Site Guide](#)
[FAQs](#)
[Mailing Lists](#)
[Contact Us](#)

Advanced

[Browser Buttons](#)
[Connotea Code](#)
[Connotea Web API](#)
[Community Pages](#)



Automating Science Protocols

Protocol: “a predefined written procedural method in the design and implementation of experiments. This should establish standards that can be adequately assessed by peer review and provide for successful replication of results by others in the field.”

(Wikipedia)

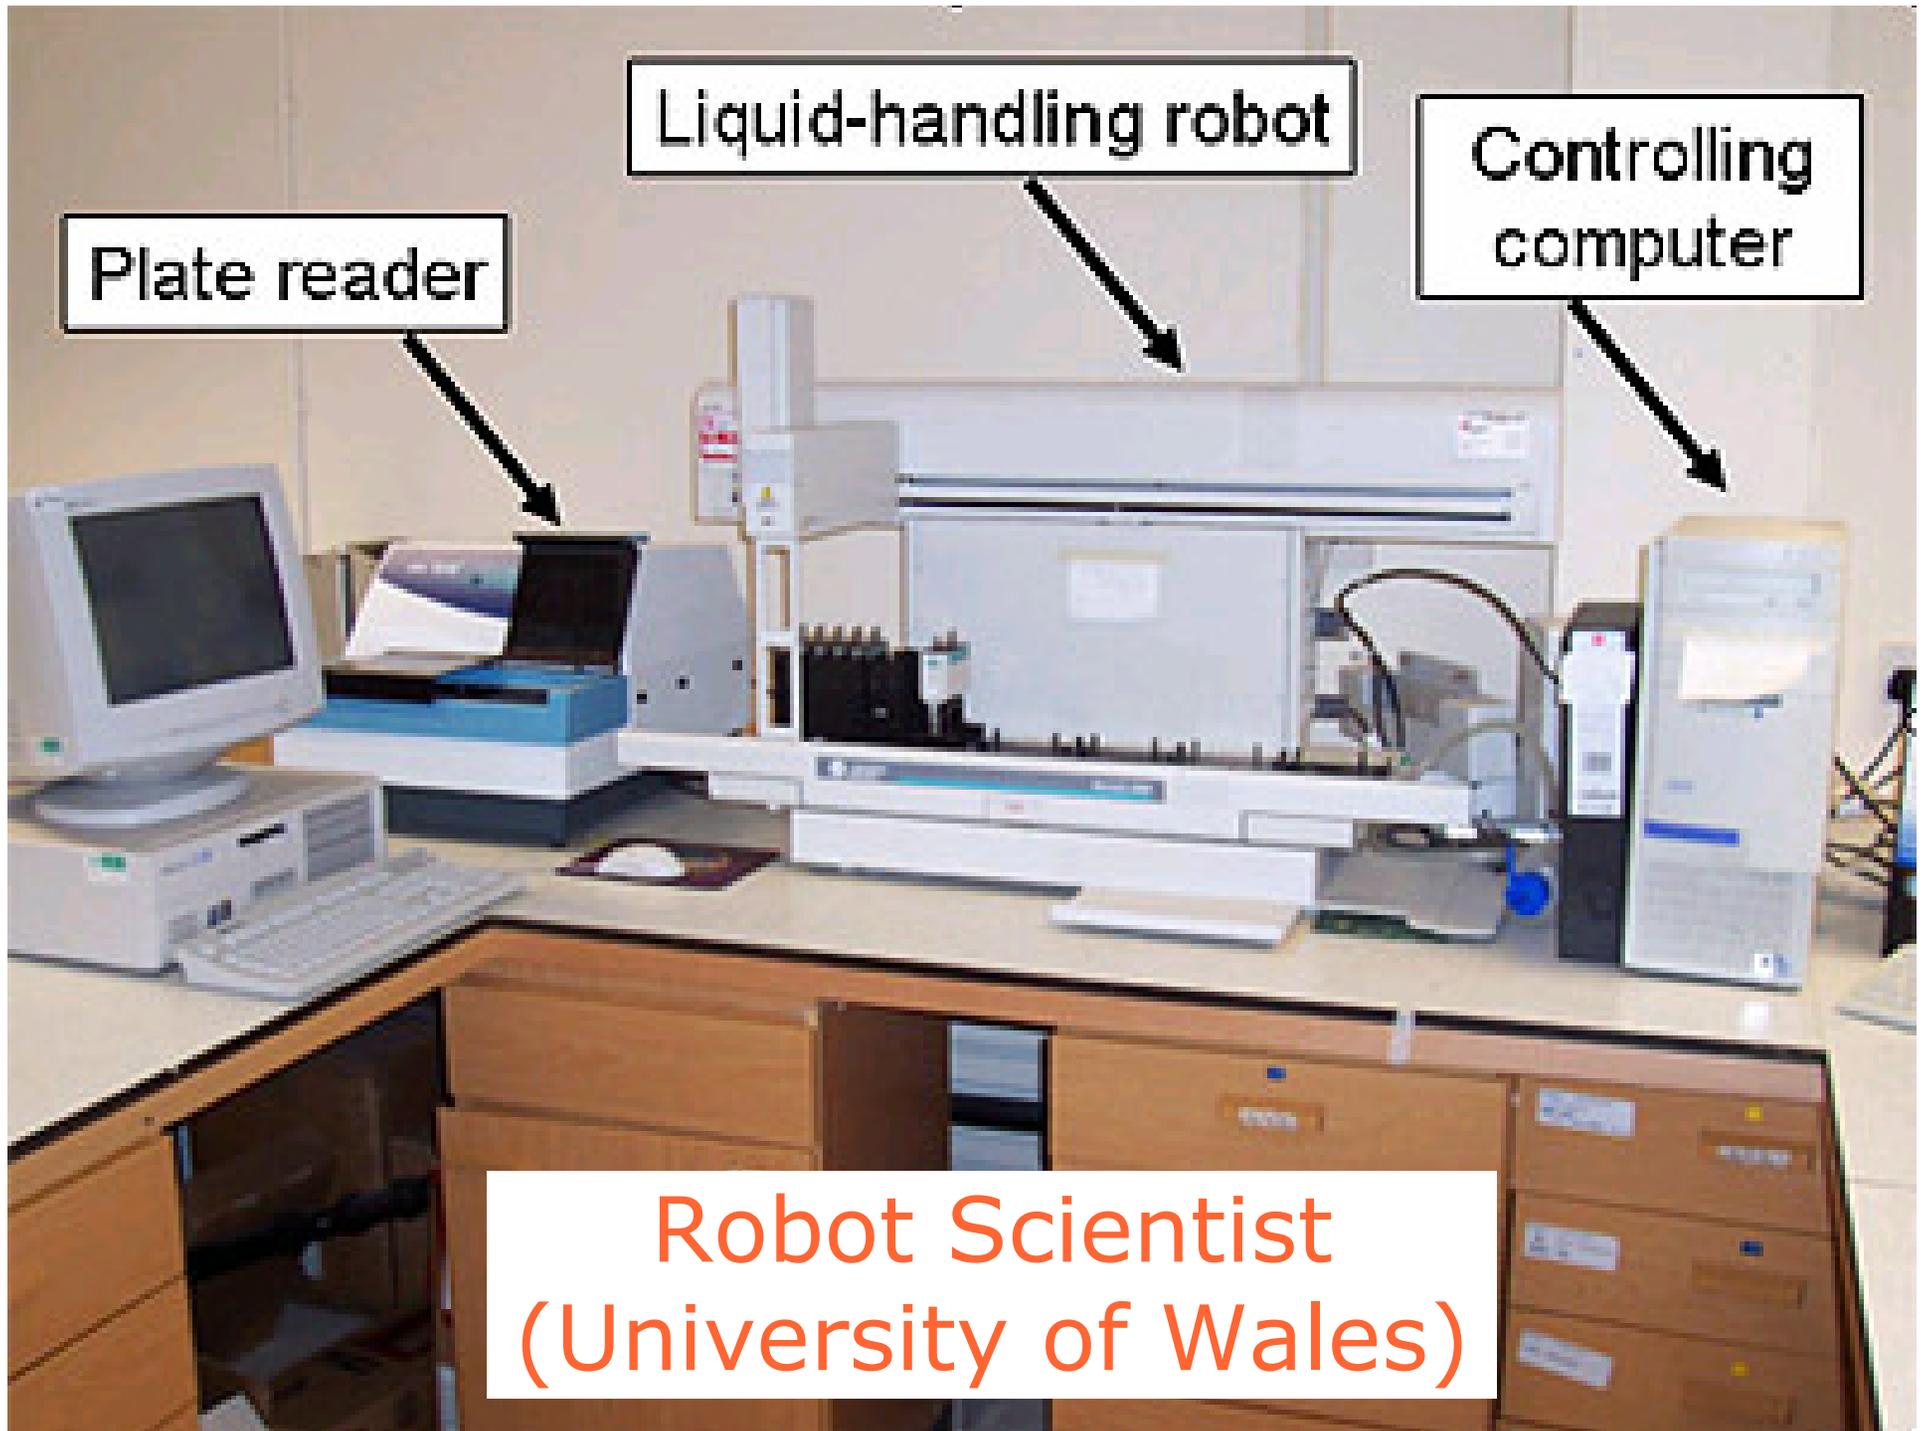


Plate reader

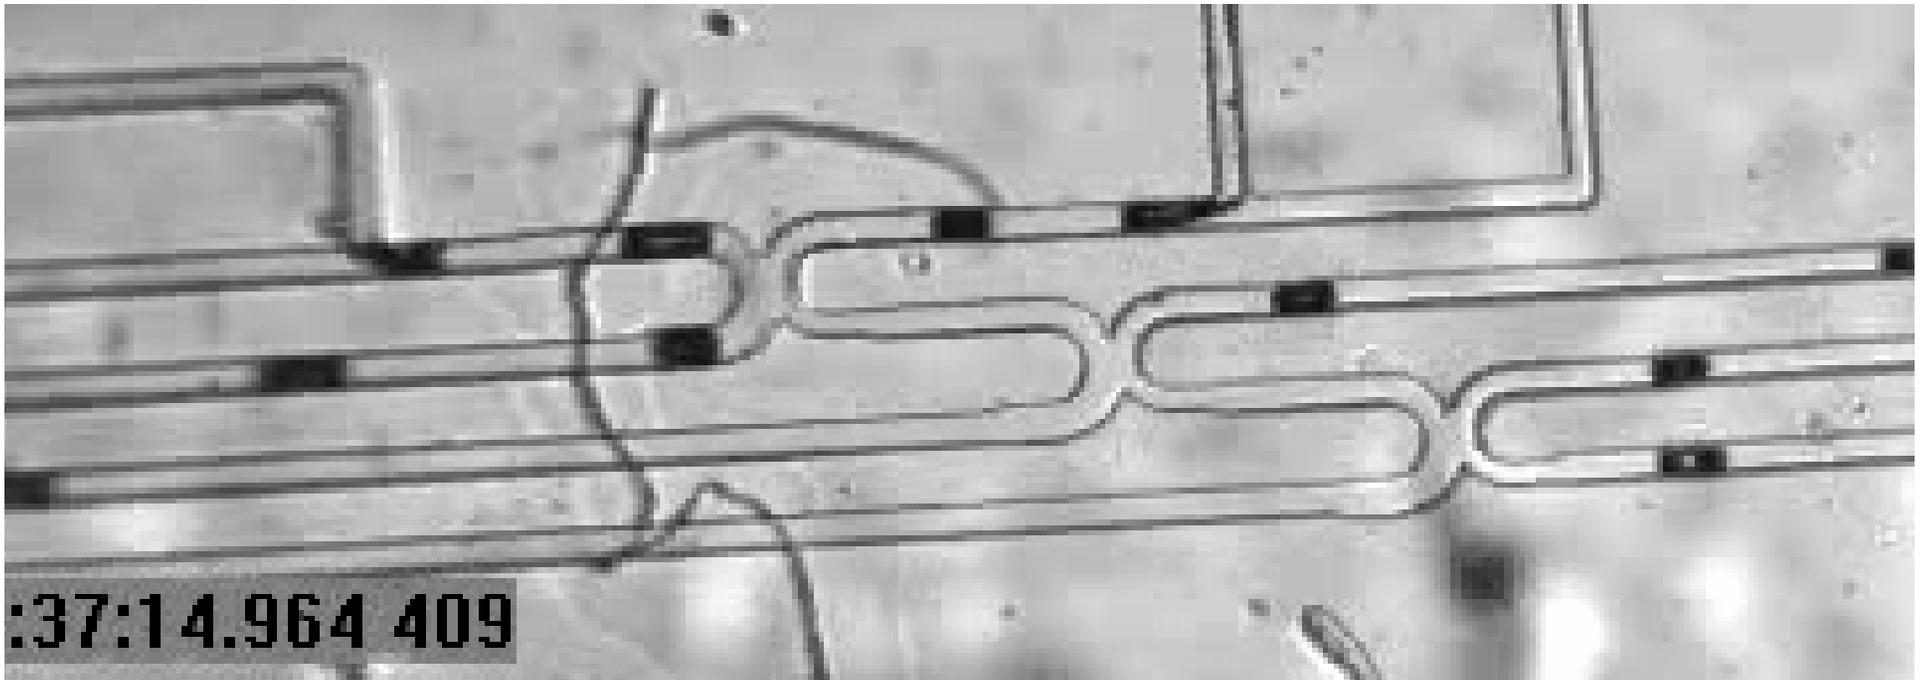
Liquid-handling robot

Controlling computer

Robot Scientist
(University of Wales)



Microfluidic Bubble Logic (Prakash and Gershenfeld)





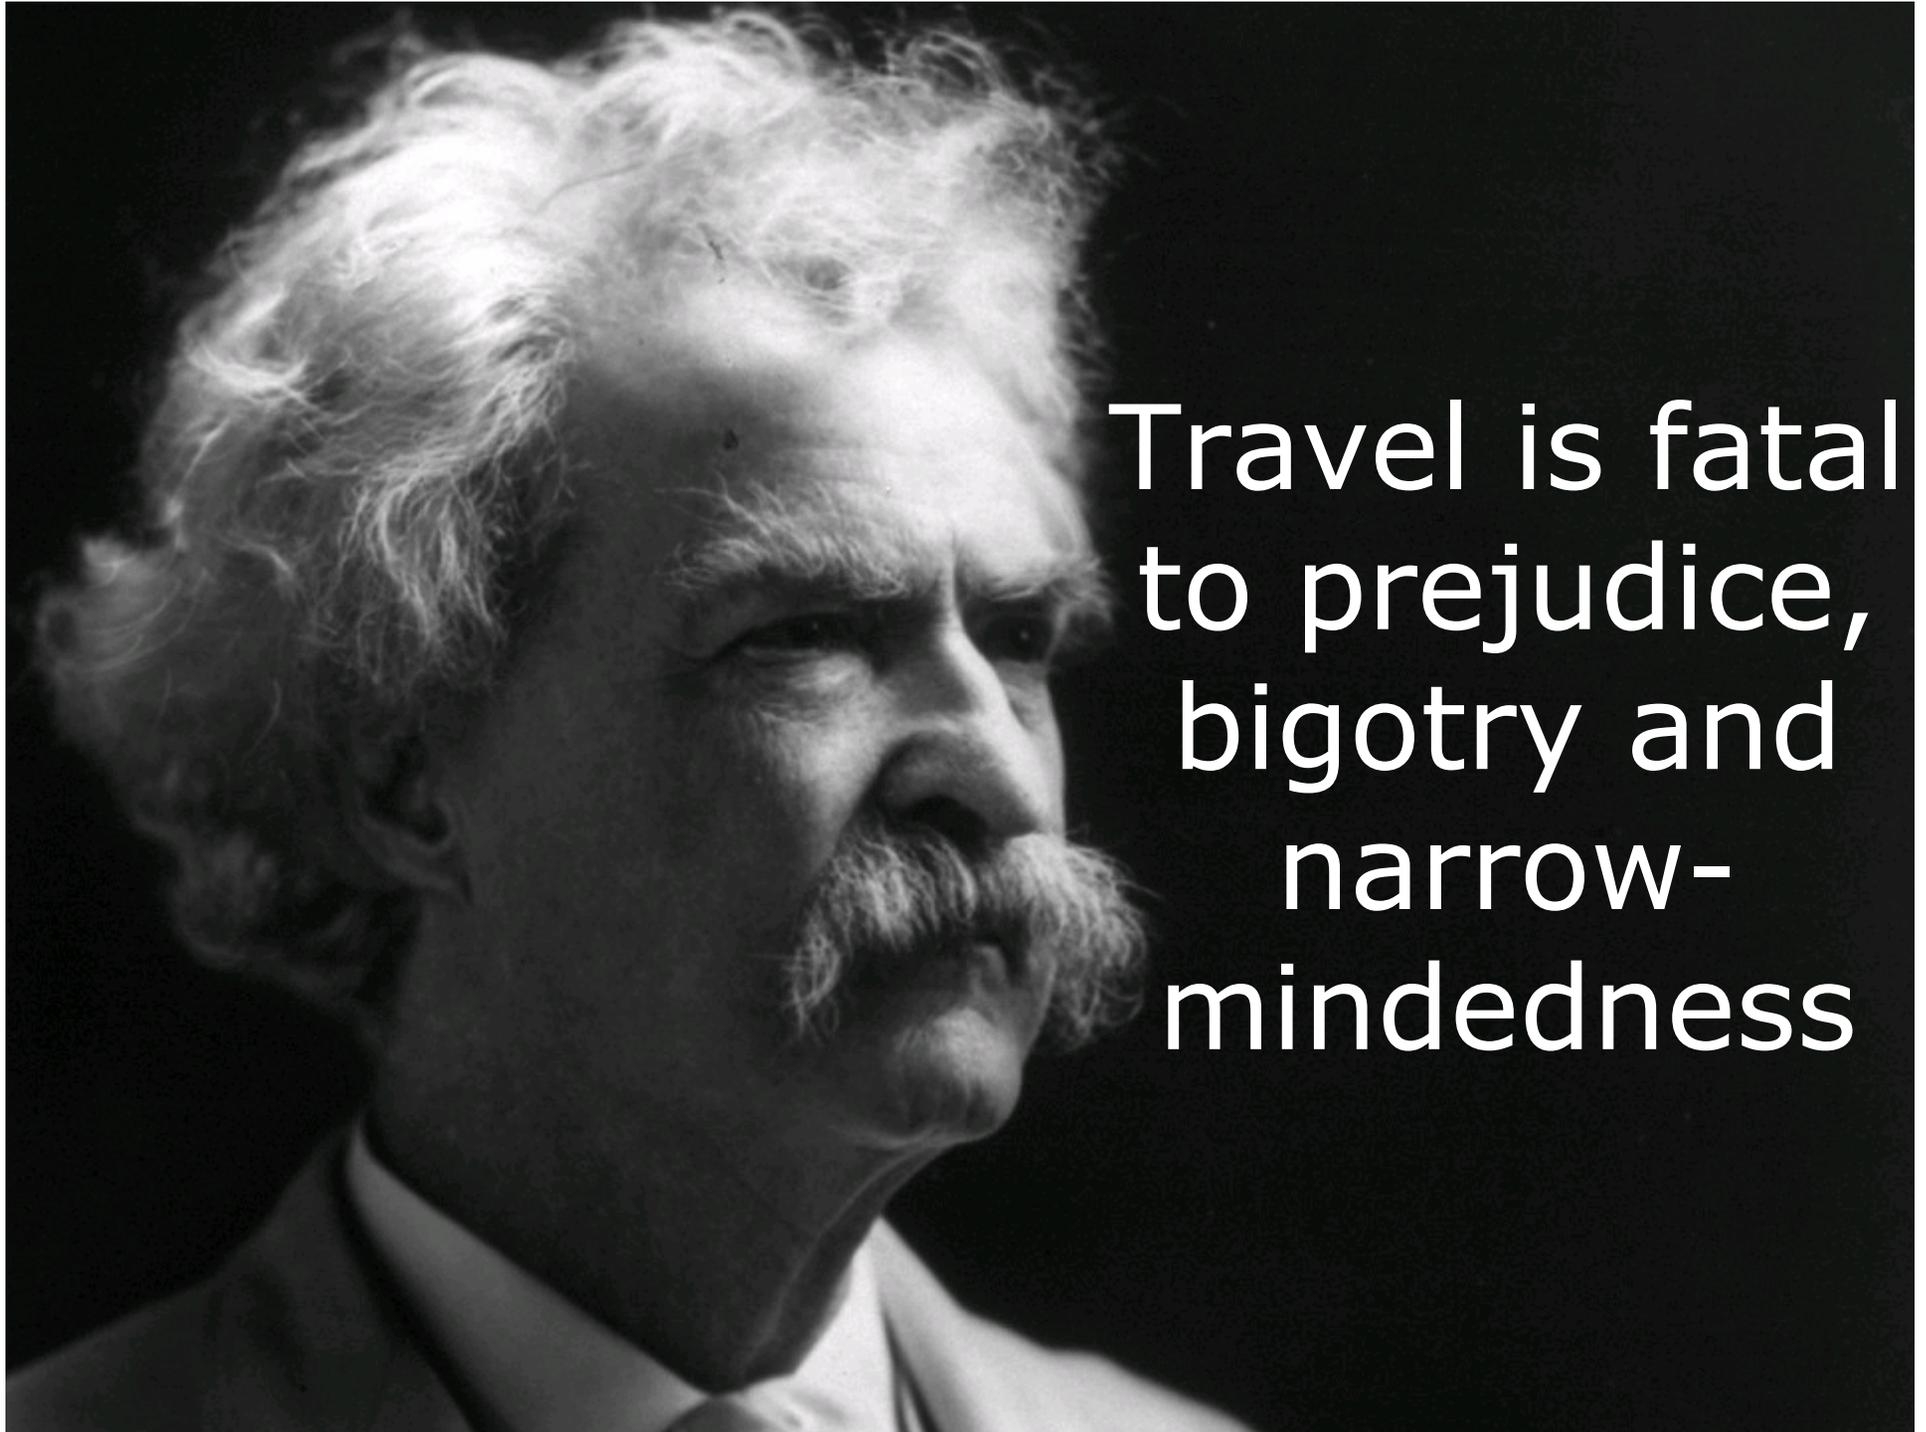
Man-Computer Symbiosis: 50 Years On

- Service oriented science
- Provenance
- Virtual organizations
- Automating science protocols



Concluding Thoughts

- Networks are necessary but not sufficient
- There is only global software
- We build communities by doing experiments
- Persistent infrastructure is needed
- Persistent workers are needed
- Showing up
- Team New Zealand?
- Where are the students?



Travel is fatal
to prejudice,
bigotry and
narrow-
mindedness